Reduced Order Modelling methodologies for proton therapy applications

Jan Willem van Galen

Research conducted as part of the Bachelor of Applied Physics program at

Delft University of Technology

Supervisor: Dr. Z. Perkó
Second examiner: Dr. ir. D. Lathouwers

Delft, July 23rd 2018

Abstract

In proton therapy, the calculation of the dose distribution is of great importance in order to find the best treatment plan. For a treatment plan with high quality, several error scenarios are investigated, in order to come up with a general plan that suits best in these scenarios. All of the scenarios have a different simulated error, e.g., in patient displacement and proton beam range. In order to cope with a large number of errors scenarios, a quick calculation method is needed. In this work, a method that used Reduced Order Modelling (ROM) and subsequently polynomial regression to overcome the computational problem is presented, as well as its results on a real head and neck cancer patient. Before error scenarios are simulated, using inverse optimisation a treatment plan is calculated for a perfectly known position and organ structure of the patient, i.e. the nominal scenario. Then, several error scenarios are simulated and for each of these scenarios, the real dose on tissue is calculated, using the treatment parameters that were previously set for the nominal scenario. The dose distributions of all these scenarios are stored in a matrix, on which a Singular Value Decomposition (SVD) is executed. A reduced order of singular values and vectors is used, together with regression models on the right singular vectors, to reconstruct the dose distribution matrix. This is done in order to create a computationally cheap method for dose distribution calculations. The right singular vectors are now a function of the used parameters: positioning errors dx, dy, dz and proton beam range error $d\rho$. The dose distribution matrix is reconstructed from the left singular vectors, first few singular values and the function for the right singular vectors. Range errors with a uniform distribution with a maximum of 3% and positioning errors with a normal distribution with a standard deviation of $3\,mm$ in each direction were simulated. For 100 error scenarios, the dose distribution matrix could be reconstructed with an average accepted error of 1% on a voxel relative to the maximum dose for 96.5% of the voxels, using a voxel-by-voxel comparison. These results were obtained for a 17^{th} order of the SVD and a 7^{th} order polynomial in the regression. On unseen test data, the acceptance was 81.0% of the voxels with an allowed error of 1%. For an allowed error of 3%, 92.5% of the voxels of this test set were accepted. These results are promising and encouraging for future research. First of all, it is recommended to examine the regression of the right singular vectors more closely, especially the regression of the higher order right singular vectors. Further subjects of interest would be to perform the same procedure that is done on the dose distribution matrix on the dose influence matrix. Also, harder-to-reach error scenarios, like organ movement and tumour deformation could be taken into account. This research was conducted at the Medical Physics and Technology section, a part of the Department of Radiation Science and Technology at Delft University of Technology. The project was part of the Bachelor program of Applied Physics at the Technical University of Delft.

Contents

1 Introduction

Worldwide, every sixth death is due to cancer and its consequences, making it the second leading cause of death, only behind cardiovascular diseases [1]. For 2016, it is estimated that 8.9 million people died from the various types of cancer. This number is estimated to grow due to expected population growth and ageing. The chance to develop cancer is 1 in 3 for both males and females, while the risk of dying from cancer is 1 in 5 [2].

1.1 Radiation therapy

Treatment of a cancer patient is most often done with (a combination of) chemotherapy, radiation therapy and surgery. For the second part, traditionally x-ray beams are used. However, this method has its disadvantages. Proton therapy is a promising alternative, because of the increased accuracy that can be reached. The use of proton therapy becomes more widespread, as several proton therapy centres are being built worldwide [3]. A representation of the difference between these two radiation methods is shown in Figure 1.

Radiating using photons leads to the undesired entrance and exit dose, as can be seen in Figure 1(a). This rest dose is a major drawback of this traditional therapy. Some improvements can be achieved, e.g., by including radiation from different angles to spread out healthy tissue damage. In Figure 1(b) it can be seen that proton has the ability to destruct a tumour more accurate, leaving surrounding tissue less damaged.

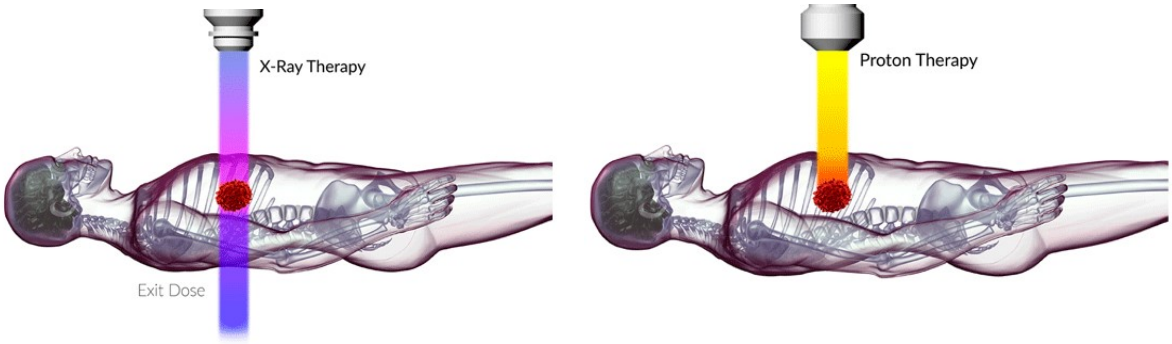

(a) With conventional x-ray radiation, the shallow tissue between the entrance and tumour receives a dose of radiation, as well as the deep tissue between tumour and exit.

(b) For proton radiation, entrance tissue still receives dose, but this dose is lower than for x-ray therapy. Additionally, when applied well, zero dose is given to deep tissue.

Figure 1: Comparison of radiation impact on the body of x-ray therapy to that of proton therapy. Image retrieved from Protom International [4].

In both cases, the process of treatment with radiating tissue is based on selected cell destruction through ionisation of an atom. This complex phenomenon will just superficially be touched in this thesis. Selected cell destruction uses a specific source to create a particle that delivers an amount of energy in a cell. The charged particle or photon will collide with an electron that is bound to an atom. This electron is released from the atom, leaving an ion. Damage on atoms leads to destruction of molecules, as part of cell structures, and eventually this leads to DNA damage. Cancerous tissue cells are less capable in repairing DNA, compared to healthy tissue, so cancerous cells will be disassembled and stop dividing more quickly.

During treatment, a patient will have to undergo several radiation sessions. When a patient would receive all dose at once, healthy tissues that get even a relatively low dose lack time to

recover from damage. When tissue is irradiated in several separated fractions, healthy cells gets time to recover, while tumourous cells miss information in DNA that gives them the capability to recover quickly.

1.2 Proton therapy and its challenge

With intensity modulated proton therapy (IMPT) it is possible to minimise the drawback of damaging surrounding tissue, by making use of the sudden sharp dose curve of protons, which can be seen in Figure 2. This characteristic Bragg-peak makes it possible to radiate the tumour shape precisely, while surrounding tissue is spared from high doses. According to the law of Lambert-Beer, given in Equation (1), the intensity I of a photon travelling through matter is an exponential function of the linear attenuation coefficient μ and depth x with the initial intensity I_0. The tissues superficial to the tumour are irradiated with a beam of higher intensity than the tumour itself.

$$I(x) = I_0 e^{-\mu x} \tag{1}$$

Protons, on the other hand, have an increased energy number of collisions just before their targeted stopping depth. The loss of energy in this Bragg peak, that can be seen in Figure 2, causes accurate local increased tissue destruction.

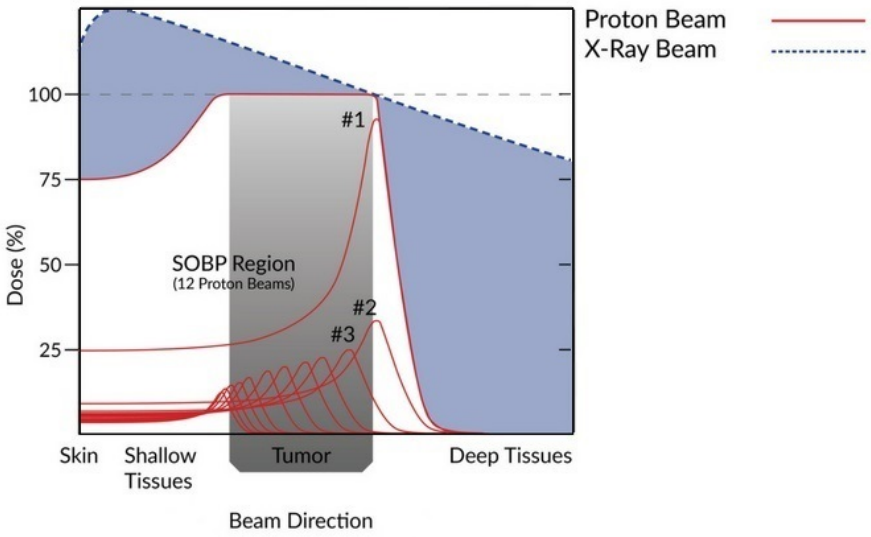

Figure 2: X-ray beams lose energy and destroy tissue following a rather slow exponential decrease. The loss in energy of proton beams is narrowly peaked and shows a sharp decrease after this Bragg-peak. Deep tissues will not receive dose. To make sure all tumourous tissue is treated well, several beams are used. A single beam would only irradiate a small local area, but several beams lead to a spread out Bragg-peak (SOBP) that covers the whole tumour.

As said, protons offer higher dose conformity compared to x-ray therapy, but this comes at a cost. The dose distribution for x-ray therapy is rather insensitive to small errors, since in case of such an error, the whole radiated area is just slightly modulated. In an error scenario, still the majority of the tumour will be irradiated properly. However, the sharper intensity edges of IMPT are more sensitive to uncertainties like patient positioning, proton beam range, internal organ motion, tumour expansion and shrinkage or anatomical changes, like weight loss [5].

The positioning error can be overcome by fixating a patient in its position, but this does not account for other errors [6]. Another way to cope this issue is using robust treatment planning. Instead of using just the nominal case, for which all uncertainties are ignored, different scenarios considering a worst case error are also taken into account. The treatment plan is calculated such that in all of these (error) scenarios, all of the constraints, i.e. treatment goals on dose distributions, are achieved as well as possible.

1.3 Aim of this project

Calculating the dose for a large number of error scenarios would increase the quality of a treatment plan, but it is computationally costly. This comes from the fact that in IMPT, each error scenario has to be recalculated, for which the intensive calculation of the dose influence matrix is needed. In this thesis, a computationally cheaper way of computing the real dose distribution under simulated errors is presented. Errors in proton beam range and patient positioning in first two and later three direction are simulated separately and combined. Using a reduced order of the singular value decomposition (SVD) of the dose distribution matrix first, and subsequently use a regression model on the right singular vectors, it is possible to reconstruct the original dose distribution matrix. This computationally cheap method is tested on a head and neck patient, by comparing to the slower accurate dose distribution calculations.

2 Theory

2.1 Treatment planning

2.1.1 Treatment requirements

For each patient that is treated, a specific treatment plan has to be calculated. The steps that have to be taken to come to such an IMPT plan are listed in Section 3.2. The requirements for such plans are given by objectives and constraints on the dose distribution. Roughly speaking should tumourous tissue absorb a higher radiation dose than healthy tissue. The requirements are drafted in dialogue with a doctor and roughly speaking state that a target dose has to be given to cancerous tissue, without damaging healthy tissues.

For the received dose, a distinction is made between the dose that irradiates a certain area and the effect that this dose has. The Relative Biological Effect (RBE) ratio compensates for the difference in biological effect that equal energies have on tissue and thus gives a measure for the loss in cell function that a radiation unit leads to. The factor is dependent on the type of radiation. IMPT uses a proton RBE relative to high-energy photons of roughly 1.1 [7]. Converting the physical dose to the RBE-weighted dose is just a matter of scaling.

2.1.2 Treatment plan calculations

The execution of treating a patient with IMPT is done using several proton beams. These protons are first accelerated with a particle accelerator, before they travel to the patient. In the IMPT plan calculations, the weight intensity that is needed for each proton beamlet is optimised. During treatment, not just a single beam is used. Instead, for each beam, a grid of bixels all shoot have beamlets with its own weight, as can be seen in Figure 3.

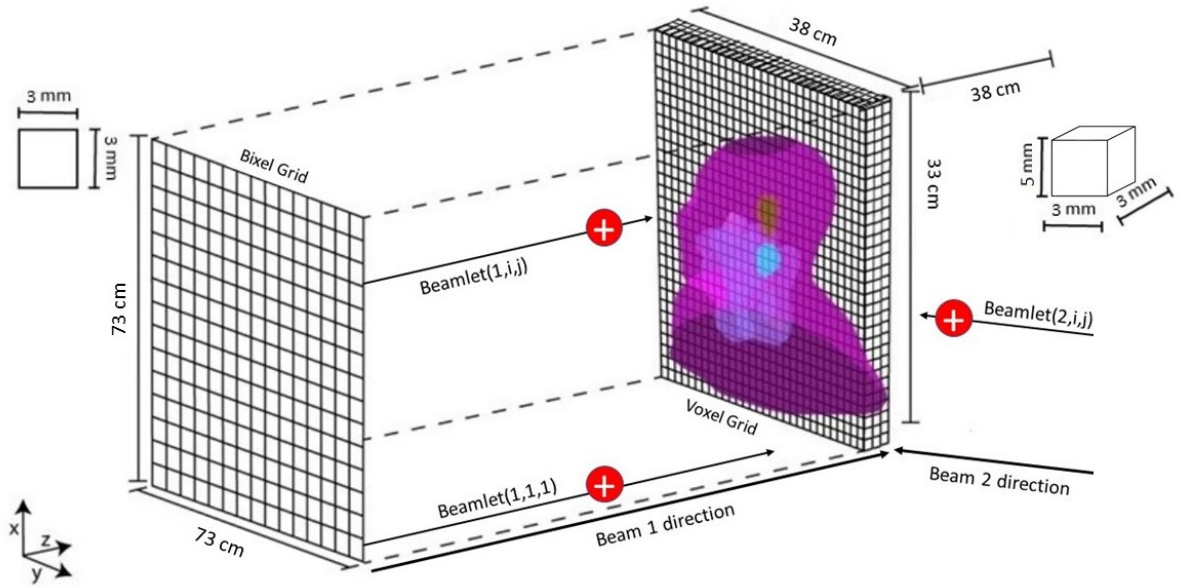

Figure 3: Visualisation of the grid of beamlet sources. On the right side only 3 slices of the image are shown, whereas in the real case there are 161 voxel in x- and y-direction and 67 in z-direction. See also Figure 5 and 6, where the CT image is given. Part of image is retrieved from Mark Arntzenius [8].

The dose due to all of the individual beams can be calculated with several algorithms. One

of them is the Monte Carlo calculation, which relies on the repeated simulation of the path of an individual proton. The method that is used during this project is a pencil beam calculation with a conventional pencil beam for IMPT, using the matRad equipment [9]. For this proton dose (D_{dose}) calculation, the fall-off is analysed in lateral (X and Y) and depth dependent (Z) components separately.

$$D_{dose} = X(x,z)Y(y,z)Z(z) \tag{2}$$

For the $X(x,z)$ and $Y(y,z)$ components, the dose fall-off is simulated with a single Gaussian in x- and y-direction respectively. In this pencil beam algorithm, the effects of materials upstream of the patient and of the air gap between source and patient are taken into account [10].

Before optimisation of the bixel weights is started, several parameters are set, including the number of beams, the angles of radiation and the bixel width. Once these are set, the influence of a certain weight on the tissue is examined bixel by bixel. By inverse optimisation, the treatment plan weight intensities are then obtained by minimising the cost function f of Equation (3), that is weighted by p_i for each of the objectives f_i, where the sum goes over all objectives i.

$$f = \sum_i p_i f_i \tag{3}$$

The objectives f_i are a sum over the errors that each voxel with index k in the objective tissue makes, so

$$f_i = \sum_k |D_k - D_{pd}|, \tag{4}$$

where D_k is the dose a voxel receives and D_{pd} is the prescribed dose for that voxel. An example for an objective f_i is the dose deviation from the mean requested dose for certain tissue. Other objective examples are overdosing of healthy tissue or underdosing of tumourous tissue. For all of these objectives, a penalty p_i is given. Optimisation of the objective function is done using the Interior Point Optimizer (IPOPT) method [11]. During the inverse optimisation, constraints on the dose are taken into account as boundary conditions. Examples of constraints are an absolute maximum or minimum dose on any voxel in a certain tissue.

Once the optimised bixel weights and the influence of the weights are known, the calculation of the total dose distribution is possible. This total dose per voxel is the sum over all the dosimetric effects of the individual beamlets.

2.1.3 Robust treatment plan calculations

One way of coping with setup and range uncertainties is to include them in the bixel weights optimisation process. With robust planning, several errors are taken into account; e.g. errors in patient positioning and organ movement. Optimising the IMPT plan is typically done using the minimax approach, in which the (maximal) error of the worst-case scenario relative to the prescription dose is minimised. The weight intensities are then inversely optimised, so that the required dose is at minimum given to a minimum fraction of the voxels; e.g. 98% of the tumourous cells receive 95% of the prescribed dose in at least 98% of the scenarios. [12] [13].

2.2 Reduced Order Modelling

In this project, factorisation of a matrix using Singular Value Decomposition (SVD) is performed. The path that is followed is visualised in Figure 4. As shown in Figure 4(a) the matrix $\boldsymbol{D}$ is decomposed in a matrix of left singular vectors $\boldsymbol{U}$, a matrix containing singular values $\boldsymbol{S}$ and a matrix containing right singular vectors $\boldsymbol{V}^T$ [14]. Shown in Figure 4(c) is the reduced

order matrix $\boldsymbol{D}_O$, which is computed from the first O left singular column vectors $\boldsymbol{U}$, the first O diagonal elements as singular values of $\boldsymbol{S}$ and the first O right singular row vectors of $\boldsymbol{V}^T$ so that

$$\boldsymbol{D}_O = \boldsymbol{U}_O * \boldsymbol{S}_O * \boldsymbol{V}_O^T. \tag{5}$$

Performing SVD and taking a reduced order is an exact decomposition, so in the limit

$$\lim_{O \to N} \boldsymbol{D}_O = \boldsymbol{D}. \tag{6}$$

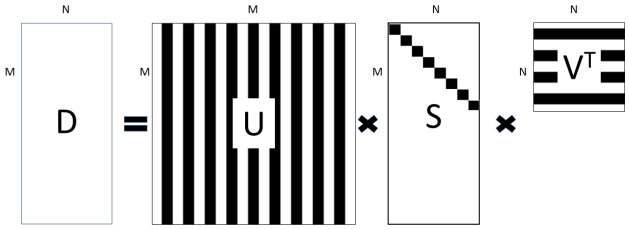

(a) Visual representation of the singular value decomposition. The information in $\boldsymbol{D}$ is represented by left singular vectors in $\boldsymbol{U}$, singular values in $\boldsymbol{S}$ and right singular vectors in $\boldsymbol{V}^T$.

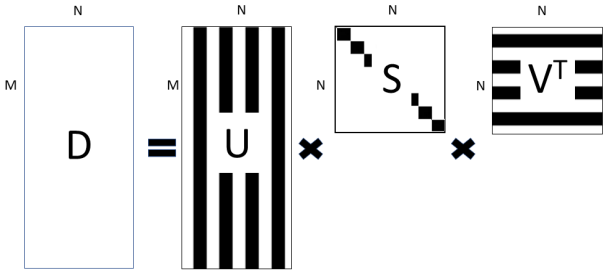

(b) A more economical version that can be used is to take the biggest possible square matrix for $\boldsymbol{S}$ and correspondingly change the size of $\boldsymbol{U}$.

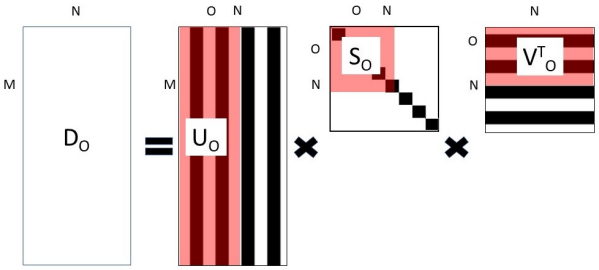

(c) Visualisation of the O'th order approximation of $\boldsymbol{D}$, where on the right-hand side of the equation only the red highlighted parts of the matrices are extracted to get as reduced order matrices $\boldsymbol{U}_O$ on the left, $\boldsymbol{S}_O$ in the middle and $\boldsymbol{V}_O^T$ on the right. Taking a reduced order will result in loss of information from $\boldsymbol{D}$, but most information is stored in the first few singular values, so keeping more terms will result in an ever-decreasing amount of information that is added to $\boldsymbol{D}$.

Figure 4: Visualisation of different types of Singular Value Decomposition of a matrix and the resulting Reduced Order Model (ROM).

3 Numerical methods and materials

3.1 Patient data

In this project, a contoured computed tomography (CT) scan of a real head and neck cancer patient from the Common Optimization for Radiation Therapy (CORT) data set is used [9] [15]. The contours label tissue of the CT image after its biological name. In this way, tumourous tissue can be separated from other organs, but most importantly from Organs At Risk (OAR). An example of such an CT image is given in Figure 5. Here only a slice is shown, for $z = 180\,mm$, whereas the whole CT image has 3 dimensions, namely x, y and z. This 3D CT image is given in Figure 6.

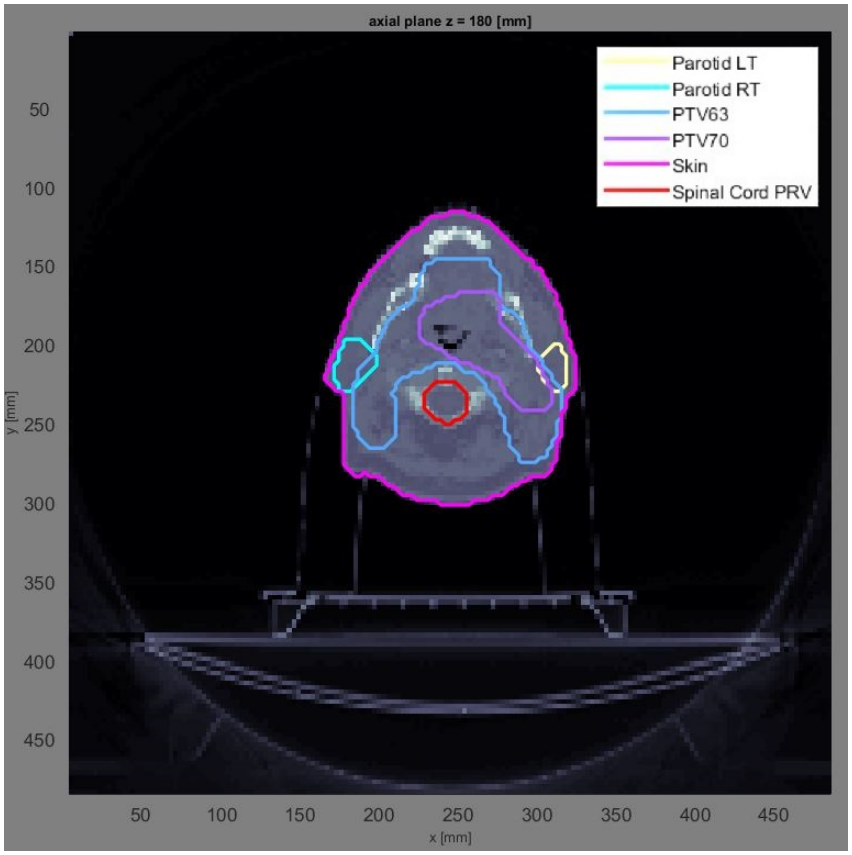

Figure 5: Slice of a head and neck CT image of a cancer patient. For the tumour, different regions are distinguished. The Gross Tumour Volume (GTV) is the region in which tumour can be seen on a CT scan. The Clinical Target Volume (CTV) indicates the region where, besides tumour visible on the CT image, microscopic cancerous cells might be found when viewed surgically. The Planning Target Volume (PTV) is the region that is planned to be irradiated, taken uncertainties of movement of the tumour into account. Only these PTVs are indicated with a desired absorbed dose of ionising radiation in units of Gray (Gy). For its vital functionality, the spinal cord is spared from high doses. The spinal cord shown in the CT image is encompassed by a Planning Risk Volume (PRV), which again takes uncertainties of movement into account, to be sure the spinal chord itself will be excluded from radiation dose. Besides, the left and right parotids are spared, because damage could lead to the dry mouth syndrome. The dose on skin is also minimised.

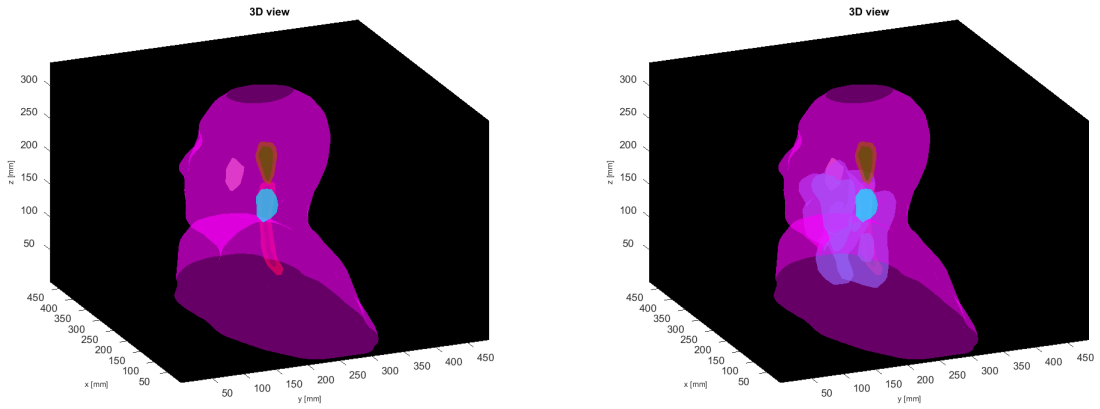

(a) Without Planning Target Volumes indicated. (b) With Planning Target Volumes indicated.

Figure 6: 3D CT images, indicating the location of the Planning Target Volumes.

3.2 Treatment plan calculations

For this project, the Matlab software matRad is used. This is open source software for three-dimensional radiation treatment planning for IMPT and other radiation therapies [9]. The steps that are taken in order to come up with a plan are briefly described here. In this section, the calculation and optimisation for only the nominal case scenario is explained.

In the first place, a 3D CT image containing patient data is loaded in the program. This image is displayed in Figure 5. With the CT comes a segmentation set that indicates tissue coordinates as well as clinical objectives and constraints (CST), as explained in Section 2.1.1. The objectives and constraints that are used are listed in Table 1. Before calculating, a plan structure (PLN) is set (See Figure 7), which includes the radiation modality and beam orientation. An automatic result is the beam geometry, i.e. the beamlet spot placements. The steering information (STF) contains settings like beam width and angle. Using these information, the dose influence matrix $\boldsymbol{D}_{i,j}$ is calculated, which gives the relation between certain bixel weights and the received voxel doses due to the weights. This dose distribution is a summation over all the lateral dose profiles from the individual beams that were radiated from the bixel spots. Then, using inverse optimisation, weights are assigned to all the bixels, in such a way that the requirements of Table 1 are as closely met as possible.

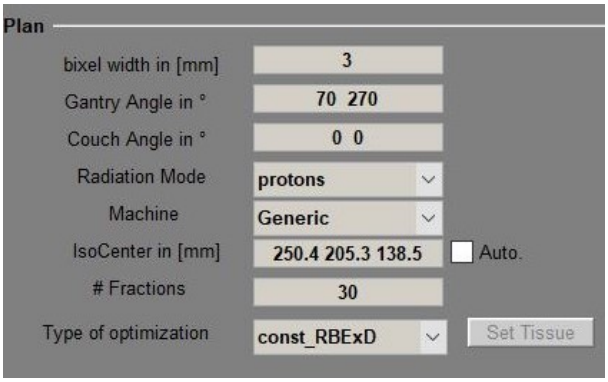

Figure 7: The plan parameters that were set for the nominal treatment plan calculation, displayed in matRads Graphical User Interface.

Table 1: Objectives and constraints (CST) for the optimisation. The volumes of interest (VOI) are assigned either as target organ or as OAR. For each VOI, several parameters on objectives or constraints are set.

Objectives & constraints

	VOI name	VOI type	priority	obj. / const.	penalty	dose
1	BRAIN_STEM_F	OAR	1	square overdosing	400	10
2	LARYNX	OAR	2	square overdosing	100	30
3	PAROTID_LT	OAR	1	square overdosing	200	30
4	PAROTID_RT	OAR	1	square overdosing	200	30
5	PTV63	TARGET	2	square deviation	1000	63
6	PTV63	TARGET	2	min dose constraint	NaN	60
7	PTV63	TARGET	2	max dose constraint	NaN	67
8	PTV70	TARGET	1	square deviation	1000	70
9	PTV70	TARGET	1	max dose constraint	NaN	75
10	PTV70	TARGET	1	min dose constraint	NaN	67
11	SKIN	OAR	3	square overdosing	800	30

3.3 Simulated errors

In a clinical case, one can come up with plenty of error scenarios that influence the received dose. The scope of this project only includes errors in patient positioning and errors in proton beam range. Other thinkable errors are organ movement, organ size change, tightened muscles and breathing.

Positioning errors are simulated with a Gaussian with a standard deviation of $3\,mm$. To create one example error scenario first, the CT images isocenter is moved in the y-direction.

After performing a positioning error in 1 direction, the range error of the proton beams is for $N = 100$ scenarios simulated with a uniform distribution between -3% and 3%. This range shift is simulated by rescaling the CTs Hounsfield Units HU. The Hounsfield scale is a measure of radiodensity, with a higher HU indicating that a unit of radiation is less able to pass the material, so that material has higher stopping power. In case of up scaled Hounsfield units, the proton beams lose their energy prior to their original stopping depth, reflecting a range error. The HU-scale is relative to water, as is shown in Equation (7). μ_{tissue}, μ_{water} and μ_{air} are the linear attenuation coefficients for tissue, water and air. By definition, $HU_{water} = 0$ and $HU_{air} = -1000$.

$$HU_{tissue} = 1000 * \frac{\mu_{tissue} - \mu_{water}}{\mu_{water} - \mu_{air}} \tag{7}$$

After the range errors, 100 error scenarios of first 2- and later 3-directional displacements are simulated, all with a standard deviation of $3\,mm$ in all directions. Lastly, 100 error scenarios with a combination of both patient displacements and range errors are investigated.

3.4 Treatment plan recalculations

3.4.1 Recalculations of dose distribution

For each of the error scenarios, the resulting dose distribution matrix is calculated. The bixel weights that were set for the nominal case will still be used, because the patient is irradiated with the beams with these beamlet weight intensities. The dose will be calculated using these non-ideal weights, without re-optimising. Both for range shifting and displacement, the dose

influence of a unit weight changes with respect to the nominal case. The new dose every voxel receives in every error scenario is stored in a matrix.

3.4.2 Representation of dose distribution

The RBE-weighted dose that each voxel gets is represented in 3 dimensional matrix form, which shape is shown in Figure 8. This dose distribution matrix $\boldsymbol{D}_M(i)$ is different for each different error scenario i. The dimensions of a single dose distribution matrix are the same as these of the CT scan: $N_x = 161$ voxels in x-direction, $N_y = 161$ voxels in y-direction and $N_z = 67$ voxels in z-direction. The total number of voxels $M = N_x * N_y * N_z = 161 * 161 * 67 \approx 1.74 * 10^6$.

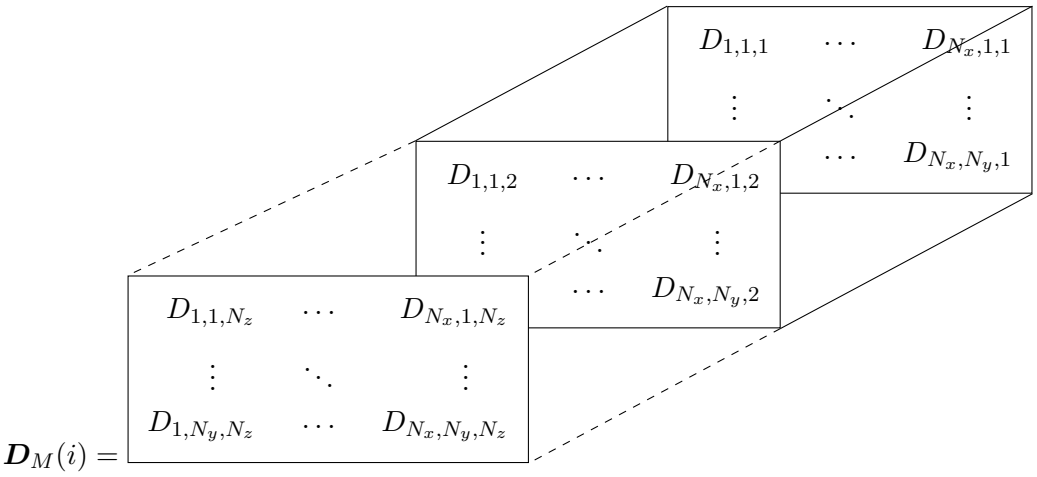

Figure 8: Dose distribution matrix, where the elements D_{N_x,N_y,N_z} are the doses that the voxel receives in error scenario i.

This matrix representation with $\boldsymbol{D}_M(i) \in \mathbb{R}^{N_x \times N_y \times N_z}$ is converted to a vector $\boldsymbol{D}_V(i)$ representation, where for one scenario all voxel doses are arranged into one vector.

$$\boldsymbol{D}_V(i) = \begin{bmatrix} D_{1,1,1} \\ \vdots \\ D_{N_x,1,1} \\ \vdots \\ D_{N_x,N_y,1} \\ D_{1,1,2} \\ \vdots \\ D_{N_x,N_y,2} \\ D_{1,1,N_x} \\ \vdots \\ D_{N_x,N_y,N_z} \end{bmatrix} \in \mathbb{R}^M \tag{8}$$

A full dose distribution matrix $\boldsymbol{D}_{fulldose}$ with the vectors $\boldsymbol{D}_V(i)$ as columns is formed. The size of this matrix depends on the number of scenarios N and the number of voxels M.

$$\boldsymbol{D}_{fulldose} = [\boldsymbol{D}_V(1) \quad \boldsymbol{D}_V(2) \quad \dots \quad \boldsymbol{D}_V(N)] \in \mathbb{R}^{M \times N} \tag{9}$$

3.5 Singular Value Decomposition of dose distributions

Before the procedure of performing a SVD is executed, the size of the matrix is reduced to make the SVD computationally cheaper. The masking is done with a cutoff dose. As visualised in Figure 9, rows for which all voxels have a dose below 0.1 Gy are removed and only the rows with at least 1 higher dose are modulated in $\boldsymbol{D}_{fulldose,chopped}$.

Then, the procedure that is described in section 2.2 is used on $\boldsymbol{D}_{fulldose,chopped}$, so that instead of the matrices of Equation (10), those of Equation (11) will be used to obtain the resultant dose distribution matrix.

$$\boldsymbol{D}_{fulldose,chopped} = \boldsymbol{U} * \boldsymbol{S} * \boldsymbol{V}^T \tag{10}$$

$$\boldsymbol{D}_{fulldose,chopped,O} = \boldsymbol{U}_O * \boldsymbol{S}_O * \boldsymbol{V}_O^T. \tag{11}$$

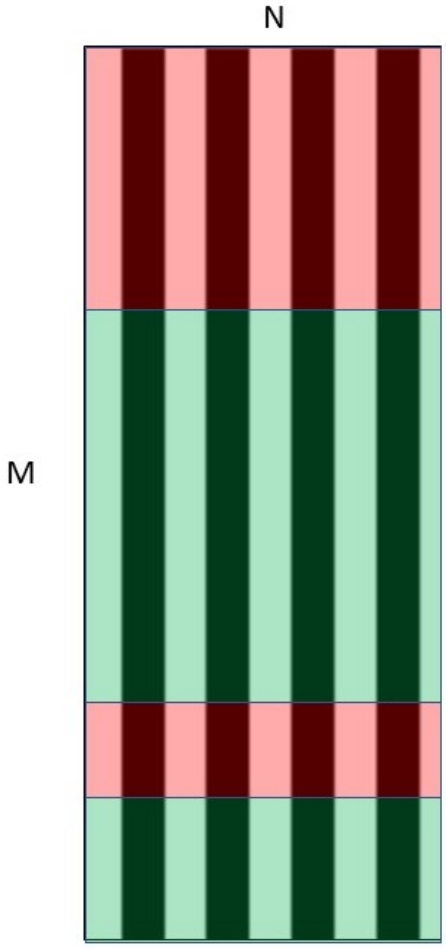

(a) $\boldsymbol{D}_{fulldose}$, where M is the number of voxels in the CT and N is the number of error scenarios. Only the green highlighted areas pass the dose mask criterion of $> 0.1\,Gy$.

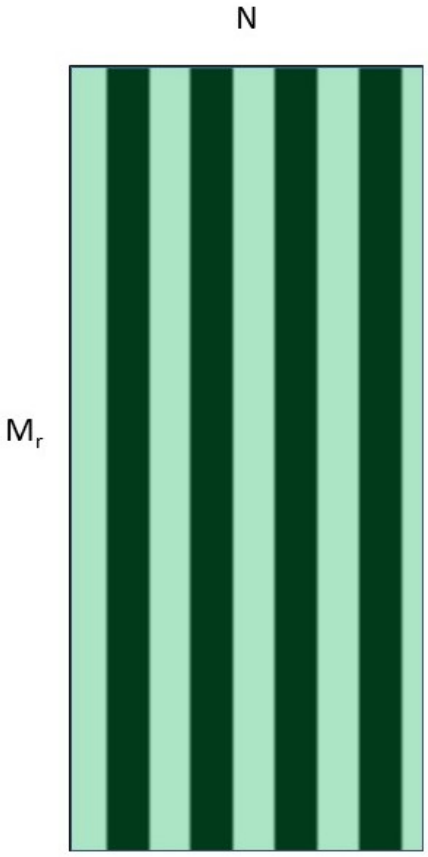

(b) Dose distribution matrix $\boldsymbol{D}_{fulldose,chopped}$ after chopping off low values, where M_r is the number of remaining voxels after chopping of low doses. N is the number of error scenarios.

Figure 9: Using a dose mask, the rows for which the RBE-weighted doses $< 0.1\,Gy$ in all N error scenarios are chopped off.

3.6 Regression of right singular vectors

The next step is to find a function

$$\boldsymbol{V}_{f,O}^T(dx, dy, dz, d\rho) \approx \boldsymbol{V}_O^T(dx, dy, dz, d\rho), \tag{12}$$

that will be used to reconstruct the right singular values by only plugging in the parameters of the error scenario. Several orders of polynomial and sinusoidal regressors will be compared to estimate $\boldsymbol{V}^T$ as close as possible. First, for 1 and 2 directional errors, MATLABs Curve Fitting Tool is used up to perform regression with 5^{th} polynomial orders and separately for 5 sinuses. For 3- and 4-directional errors, a regression model with up to 3^{rd} order polynomial crossterms and up to 7^{th} order polynomial separate terms is used. In all cases, most desirable is to have as less parameters as possible, while preserving sufficient accuracy.

Using the estimated function, calculation of $\boldsymbol{V}^T$ can be done quickly, while the matrix multiplication of Equation (13) is also computationally cheap. Together this paves the path for quick dose distribution calculations. The parameters can be filled in and the resultant dose is calculated, using the matrices $\boldsymbol{U}$ and $\boldsymbol{S}$, which were calculated from the SVD of the nominal case.

$$\boldsymbol{D}_{fulldose,chopped,f,O}(dx, dy, dz, d\rho) = \boldsymbol{U}_O * \boldsymbol{S}_O * \boldsymbol{V}_{f,O}^T(dx, dy, dz, d\rho), \tag{13}$$

where

$$\boldsymbol{D}_{fulldose,chopped,O} \approx \boldsymbol{D}_{fulldose,chopped,f,O}(dx, dy, dz, d\rho). \tag{14}$$

$\boldsymbol{D}_{fulldose,chopped,O}$ is the reduced order, full dose distribution matrix, according to how it is visualised in Figure 4(c). The reduced order, regressed, full dose distribution matrix $\boldsymbol{D}_{fulldose,chopped,f,O}$ is easy to compute, when the input variables are known.

3.7 Reconstructing the dose distribution matrices and evaluating the accuracy of the Reduced Order Model and regression

As mentioned in Section 2.2, the SVD is exact in the limit $O \to N$. However, taking only the first O order will result in some loss of information, unless the left out singular values are 0. To analyse the effect of reducing the order, $\boldsymbol{D}_{fulldose,chopped,O}$ is compared to $\boldsymbol{D}_{fulldose,chopped}$ with a voxel-by-voxel comparison using Equation (15), where i and j represent the position of the voxel in the dose distribution matrix. The difference between the voxels is divided by the maximum dose D_{max} of $\boldsymbol{D}_{fulldose,chopped}$. The fraction of the voxels that is within acceptable margin, having a percentage error in dose ϵ_{dose} is calculated. First this is done for an average over N scenarios, then also for the worst case error scenarios, where the range error $d\rho$ or the displacement error $ds = \sqrt{dx^2 + dy^2 + dz^2}$ is biggest.

$$\epsilon_{dose}(i,j) = \frac{\boldsymbol{D}_{fulldose,chopped,O}(i,j) - \boldsymbol{D}_{fulldose,chopped}(i,j)}{D_{max}} \tag{15}$$

To see how well the regression on the right singular vectors works, $\boldsymbol{V}_{f,O}^T$ is compared to $\boldsymbol{V}_O^T$. In order to do this comparison, the coefficients are scattered and the estimated function is plotted. For quantitative insight, the residuals are calculated.

For comparing the regression of the training set to the data of the test set, its right singular vectors are calculated using Equation (16)-(18), where the pseudo-inverse is used.

$$\boldsymbol{D}_{fulldose,train,chopped,f,O} \approx \boldsymbol{U}_{train,O} * \boldsymbol{S}_{train,O} * \boldsymbol{V}_{train,f,O}^T \tag{16}$$

The following step is sort of a forced SVD, because $\boldsymbol{U}$ and $\boldsymbol{S}$ are not chosen freely. Still, in the limit $O \to N$ this is an exact decomposition.

$$\boldsymbol{D}_{fulldose,test} \approx \boldsymbol{U}_{train,O} * \boldsymbol{S}_{train,O} * \boldsymbol{V}^T_{test,f,O} \tag{17}$$

And to extract the new right singular vectors on the right side only

$$(\boldsymbol{U}_{train,O} * \boldsymbol{S}_{train,O})^{-1} * \boldsymbol{D}_{fulldose,test} \approx \boldsymbol{V}^T_{test,f,O}. \tag{18}$$

The resultant right singular vectors for the test set are scattered in the same image as described for the train set.

To see how well the regression works on the dose distribution matrix, $\boldsymbol{D}_{fulldose,chopped,O}$ is compared to $\boldsymbol{D}_{fulldose,chopped,f,O}$, again using the voxel-by-voxel comparison.

To see how well the combination of the ROM and the regression works, $\boldsymbol{D}_{fulldose,chopped,f,O}$ is compared to $\boldsymbol{D}_{fulldose,chopped}$ using voxel-by-voxel comparison. This is also done for unseen test data, where $\boldsymbol{D}_{fulldose,chopped,f,O}$ is compared to $\boldsymbol{D}_{fulldose,chopped,test}$. For the test data, still the dose mask is used, and only the rows that were included for training data are examined.

4 Results and Discussion

4.1 Nominal treatment plan results

To start calculations, a treatment plan is calculated for the scenario in which the exact position of the patient is known and fixed. This is called the nominal scenario. When the steps of Section 3.2 are followed, the dose that every voxel gets is calculated with this treatment plan. In Figure 10, the resulting dose for each voxel is shown. The real image is a 3D image, so here only a single slice is shown.

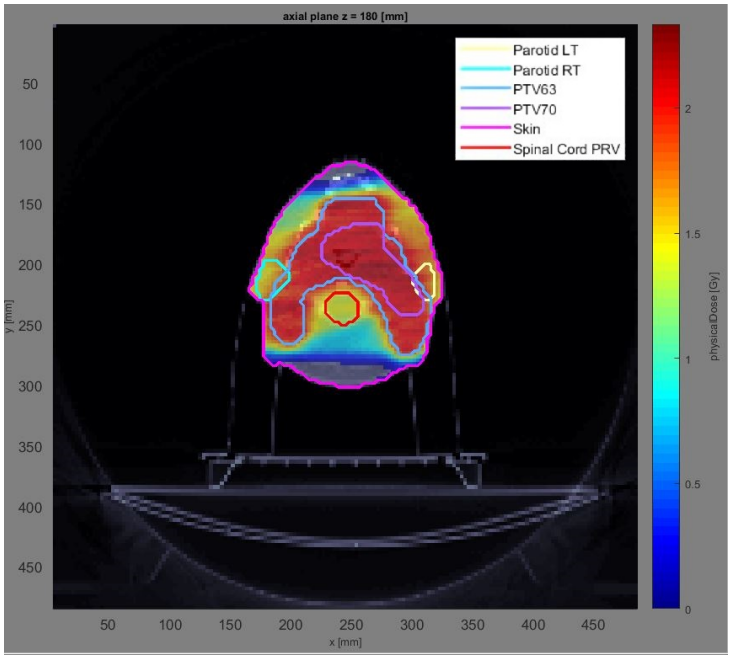

Figure 10: Physical radiation dose per fraction is plotted in the CT image of Figure 5. It can be seen that the PTV receives a high dose compared to the surrounding tissue. Even at the PTV boundaries, the dose is relatively high. The spinal cord PRV, skin and parotids are spared from high doses.

The resulting dose can be analysed using a Dose Volume Histogram (DVH), which is done in Figure 11(a). With a DVH, the fraction of an organ that receives at least a certain dose can be obtained. From a DVH, the exact distribution for each voxel can not be read, so it works complementary to Figure 10. In Figure 11(a) it can be seen that the dose falloff for both PTVs is rather steep between 2.2 and 2.5 Gy. This means that almost all of the tumour volume receives well-nigh the same dose. This outcome is preferred over a dose with a higher deviation. The received dose that is shown here is the RBE-weighted dose.

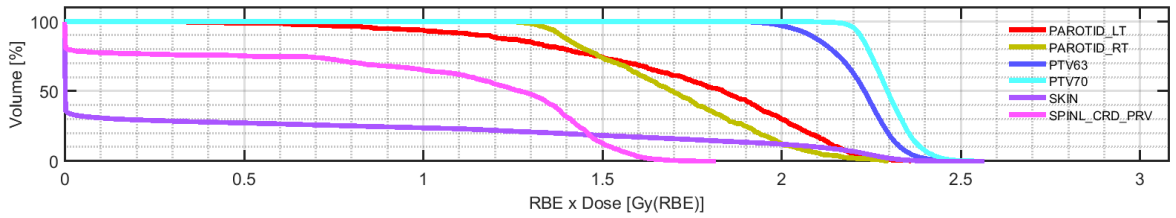

(a) DVH of the nominal case scenario. The mean dose for the PTV63 is $2.2\,Gy$ and the mean dose for the PTV70 is $2.3\,Gy$. The mean dose on the spinal cord is $0.091\,Gy$. This is an acceptable plan, which still could be improved, e.g. by increasing the number of beams and optimising the radiation angles.

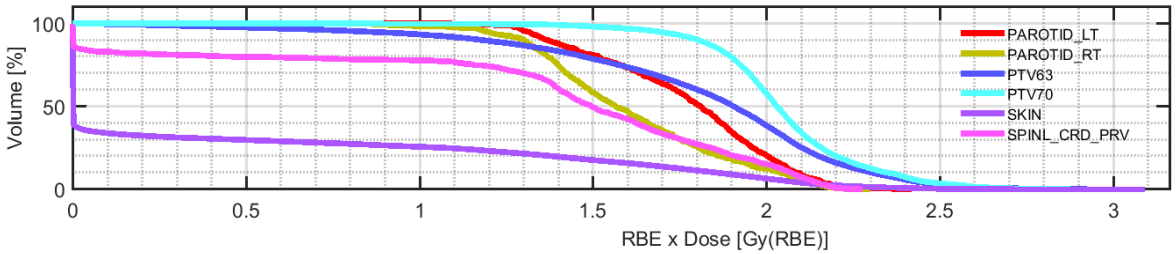

(b) DVH of the error case scenario, where $dy = 14\,mm$. The mean dose for the PTV63 is $1.8\,Gy$, for the PTV70 it is $2.0\,Gy$. The mean dose on the spinal cord is $0.4\,Gy$.

Figure 11: DVH of the nominal scenario and of an error scenario. Shifting the patient by $14\,mm$ has a significant influence on the quality of the dose distribution that results from the nominal treatment plan. In the error scenario, the PTV70 that includes the tumour should receive a dose of $23.3\,Gy$ per fraction, but is on average underdosed by $0.3\,Gy$ per fraction and the spinal cord PRV, that should be protected, is on average overdosed by $0.3\,Gy$. This is the result of not taking this error scenario into account in the optimisation of the plan.

4.2 Demonstrating the sensitivity of IMPT treatments with an error in y-direction

To start the calculation of dose distributions in error scenarios, an error in the placement of the patient is simulated. In this case, the patient is moved by $-14\,mm$ in y-direction. The dose is recalculated, following the steps of Section 3.4.1. The resultant dose distribution is shown in Figure 12(b). The difference between the nominal and error scenario is shown in Figure 12. What catches the eye is the blue and red horizontal bar in Figure 12(c), indicating respectively a decreased and increased received dose. This difference becomes clearer in the DVH (see Figure 11(b)). In this error scenario, the plan would not be sufficiently adequate. This demonstrates that for calculation of an IMPT plan, it is of importance to know these inadequacies beforehand, in order to optimise the bixel weights also for such error scenarios.

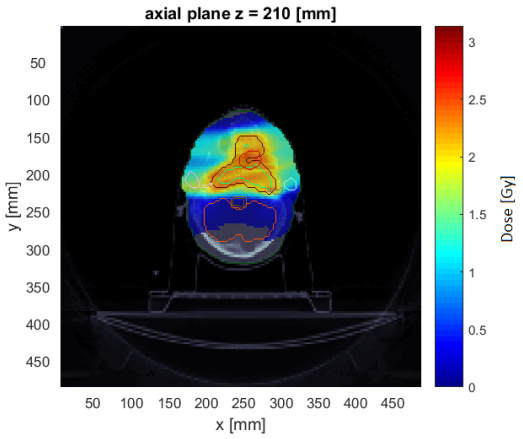

(a) Resulting physical dose distribution in the nominal position for the patient.

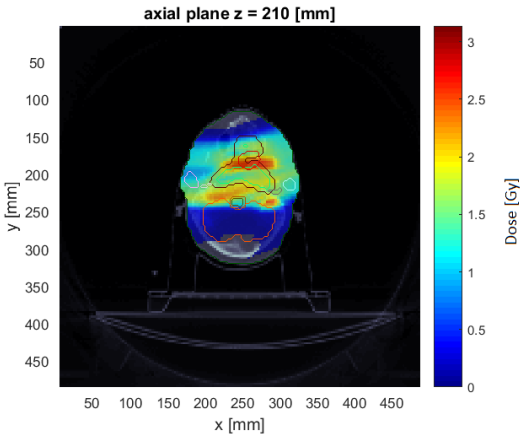

(b) Resulting physical dose distribution in the error scenario, for which the patient is shifted by $dy = -14mm$. In this figure, one can already see that part of the dose missed its target, while the spinal cord receives too much dose.

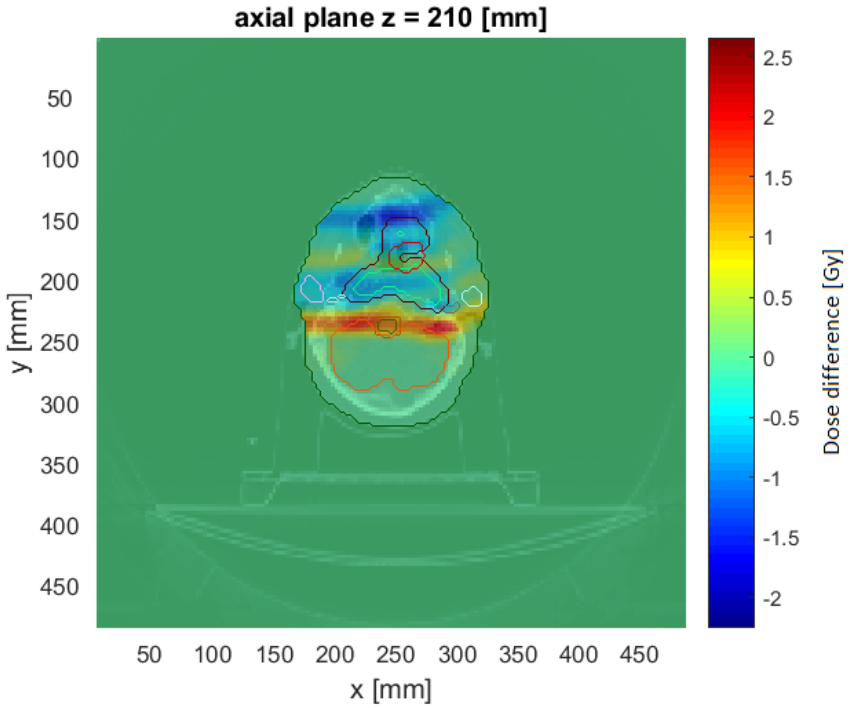

(c) Difference between the dose distributions of the nominal case and the displaced case. Like in the DVH of Figure 11(b), it can be seen that the spinal cord PRV is overdosed and the PTV at around $y = 150\,mm$ is underdosed relative to the nominal plan.

Figure 12: The resulting dose distributions in the nominal and a displacement scenario compared in a CT image slice.

4.3 Range errors

The problem of having an inadequate treatment plan in an error scenario, like in Section 4.2, can be overcome. But to know the influence on error scenarios, the dose distribution in these error scenarios should be calculated before optimising the bixel weights.

In this section, range errors will be analysed. With the procedure mentioned in Section 3.5, the rows for which every scenario has a low dose are removed. This results in a decrease from $M \approx 1.7 * 10^6$ rows to a remaining $M_r \approx 7.4 * 10^4$ rows, which makes the SVD computationally cheaper.

4.3.1 Order approximation of singular values

The dose distribution matrix $\boldsymbol{D}$, resulting from 100 range errors, has size $7 * 10^4$ voxels by 100 range error scenarios and is decomposed using Equation (5). The singular values are plotted in Figure 13. It is important to see the decrease in these singular values, knowing that the first principal component contains more information than the second, and the later ones are of ever decreasing importance. The first 5 principal components show a decrease of 4 orders in magnitude. Because of this decreasing contribution, from the next section on only the first O principal components will be taken into account, based on the accuracy an order gives.

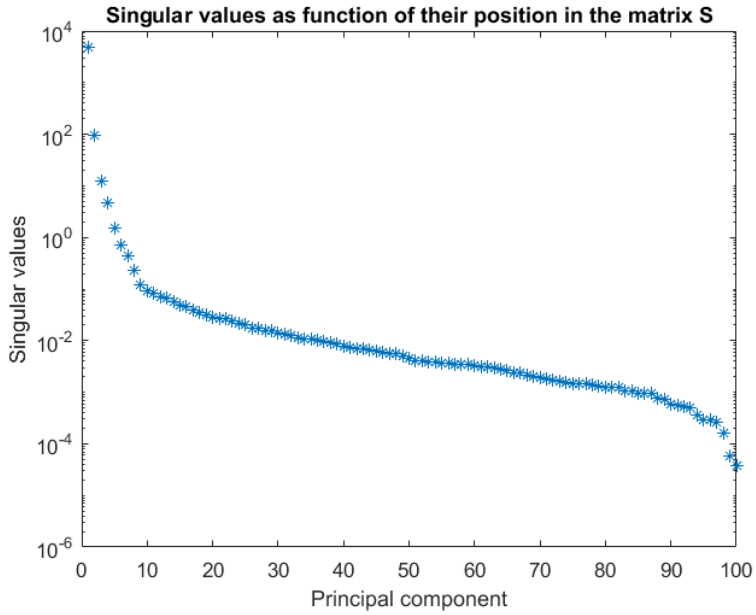

Figure 13: Singular value decrease for the range error scenarios, with a range error between -3% and 3%, where a uniform distribution is used to simulate errors.

4.3.2 Reconstructing the dose distribution matrix after using a reduced order of the SVD

The chopped dose distribution matrix is reconstructed using Equation (11). This reconstructed matrix is compared to the original dose distribution matrix, and the accuracy of the reconstruction is analysed. In Figure 14(a), the number of voxels that are reconstructed with a certain error are plotted for several orders, where the accuracy is averaged over the 100 scenarios. In Figure 14(b), the accuracy is plotted for the worst case scenario, where the range error was biggest.

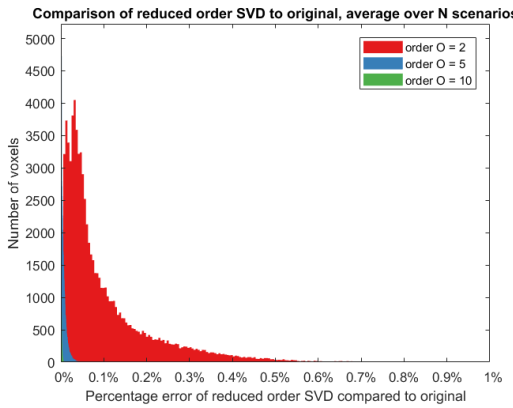

(a) A comparison between the dose distribution for several O^{th} orders of the SVD and the calculated dose distribution, as an average over the 100 scenarios.

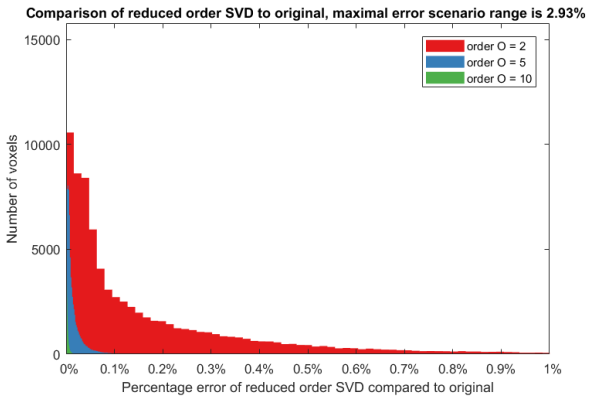

(b) A comparison between the dose distribution of several orders of the SVD, and the calculated dose distribution, for the worst case scenario, where the range error was 2.93%.

Figure 14: Number of voxels having a certain error plotted for several orders, both for the average and a worst case scenario.

Using the error histogram (Figure 14), we can determine the order O that is required for a certain accuracy. The voxels on the left side of an allowed percentage error are the voxels that are calculated to have at least that chosen accuracy. From Figure 15, it can be seen that in order to recompute 99.9% of the dose distribution matrix voxels with an accuracy of 0.5%, only 4 orders of the SVD are needed. For a required fraction of voxels with a margin of 0.1% in 99.9% of the voxels, an order O of 5 is needed, as can be seen in . This most strict requirements, i.e. an accuracy of 0.1% in 99.9% of the voxels, are set. Hence, in the continuation, only the first 5 principal components and first 5 eigenvectors in $\boldsymbol{V}^T$ are taken into account, so that

$$\boldsymbol{D}_{fulldose,chopped,5} = \boldsymbol{U}_5 * \boldsymbol{S}_5 * \boldsymbol{V}_5^T. \tag{19}$$

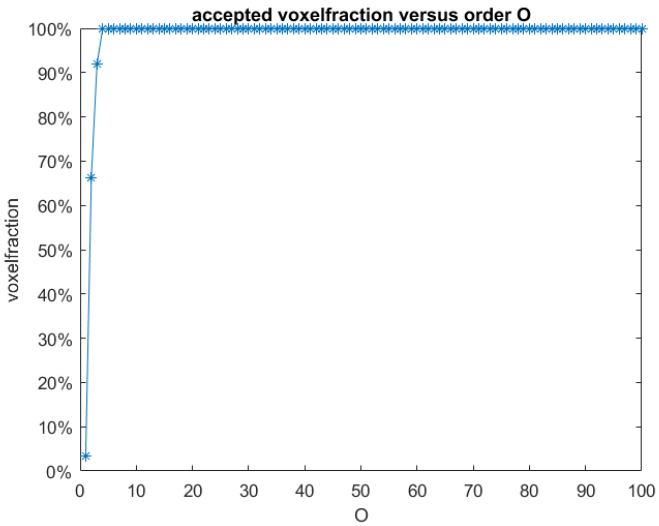

Figure 15: Fraction of voxels having a smaller error than 0.1% per voxel, averaged over 100 scenarios, as a function of the number of orders used. $O = 4$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 5$ is required.

4.3.3 Regression

On the right singular vectors a regression is performed with polynomials up to the 5^{th} order, which can be seen on the left-hand side in Figure 16. This regressions are also compared to regressions with 5 sinuses, on the right-hand side of that figure. For the first three rows, the polynomials perform better. On the latter ones, the sinuses perform better, but as these rows are less important, the calculation in this section is continued with polynomial regression of the 5^{th} order.

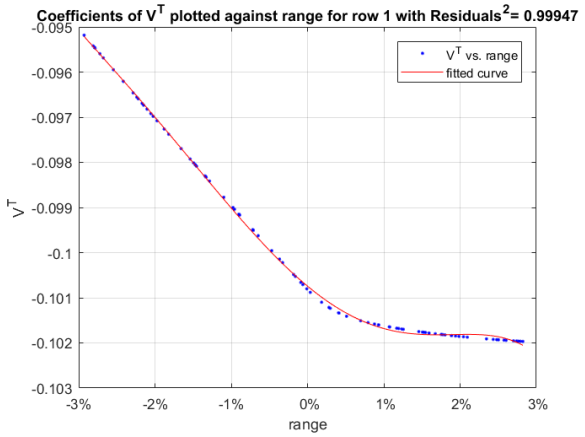

(a) Coefficients of row 1, as a function of the randomly generated range, with the regression function of 5 polynomials.

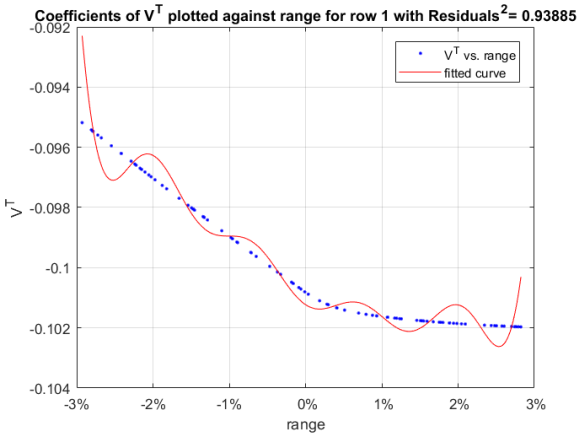

(b) Same as the figure on the left, but regression is now performed with 5 sinuses.

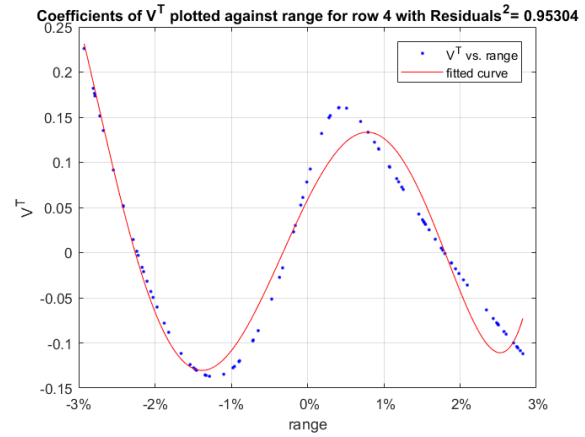

(c) Coefficients of row 4, as a function of the randomly generated range, with the regression function of 5 polynomials.

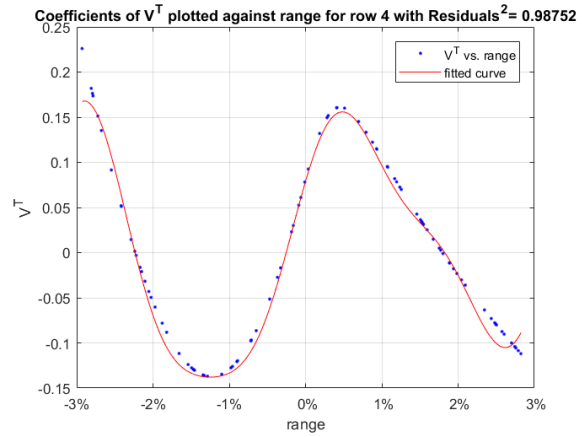

(d) Same as the figure on the left, but regression is now performed with 5 sinuses.

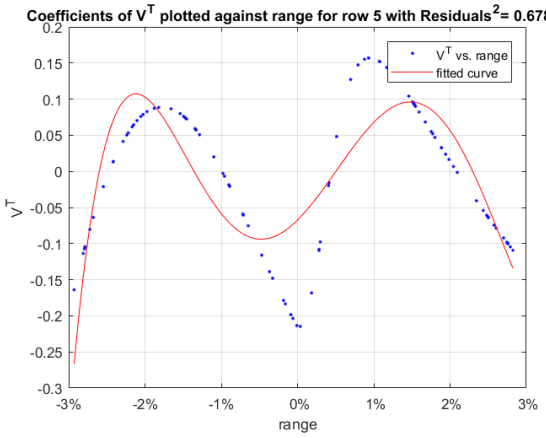

(e) Coefficients of row 5, as a function of the randomly generated range, with the regression function of 5 polynomials.

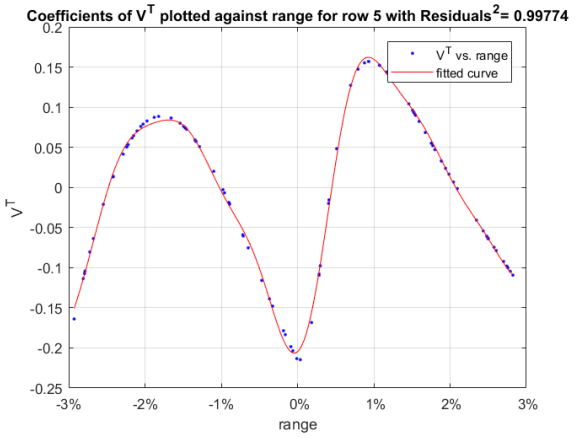

(f) Same as the figure on the left, but regression is now performed with 5 sinuses.

Figure 16: Plots of the coefficients in $\boldsymbol{V}^T$ versus range errors, with the fitted regression curves and corresponding residuals.

4.3.4 Performance on train and test data

After performing regression, the resulting matrix is compared to the matrix before the regression. Also, the right singular vectors with regression are used for reconstruction of the original dose distribution matrix and this matrix is compared voxel-by-voxel. As shown in Figure 17(b), 92.7% of the voxels have a value within margins of 0.1%.

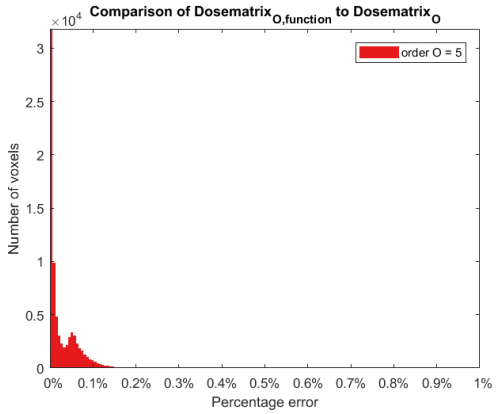

(a) Comparing the reduced order matrix with regression to the reduced order matrix results in 93.0% acceptance of the voxels, allowing 0.1% difference. All voxels have an error below 1%.

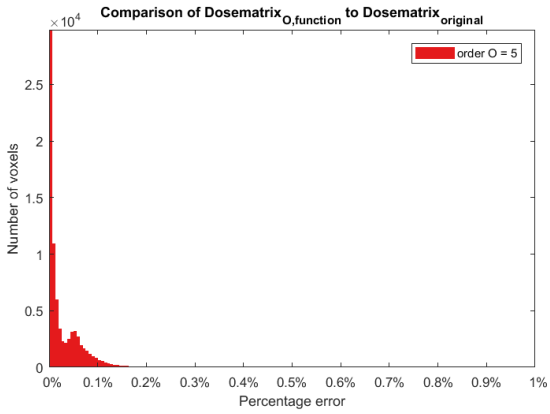

(b) Comparing the reduced order matrix with regression to the original matrix results in 92.7% acceptance of the voxels, allowing 0.1% difference. All voxels have an error below 1%.

Figure 17: Comparison between the reconstructed matrices and the original matrix, for the average over all scenarios. Regression is done with a 5^{th} order polynomial and $O = 5$.

Reduced Order Modelling and subsequently regression on the right singular vectors works well on the training set. Now using the procedure of Section 3.7, the coefficients for unseen test data are obtained and investigated. For several rows, the coefficients are compared with the training data and corresponding regression curves in Figure 18.

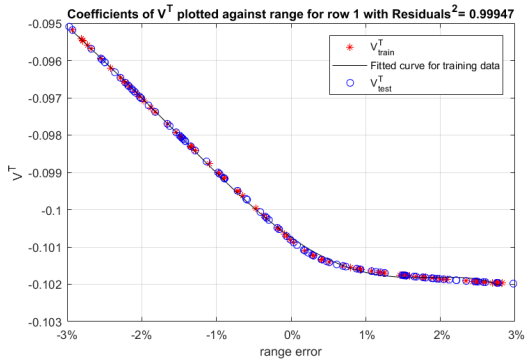

(a) First row of $\boldsymbol{V}_{train}^T$ compared to $\boldsymbol{V}_{test}^T$, plotted with the regression function.

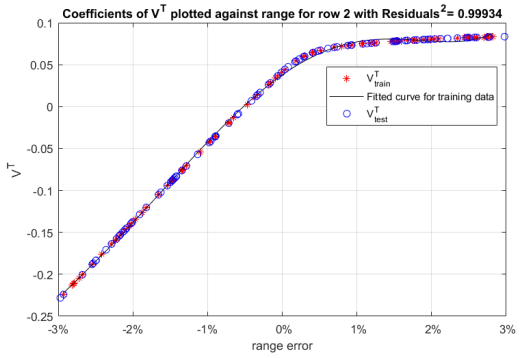

(b) Second row of $\boldsymbol{V}_{train}^T$ compared to $\boldsymbol{V}_{test}^T$, plotted with the regression function.

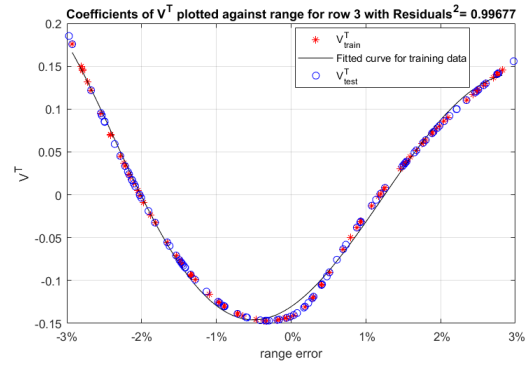

(c) Third row of $\boldsymbol{V}_{train}^T$ compared to $\boldsymbol{V}_{test}^T$, plotted with the regression function.

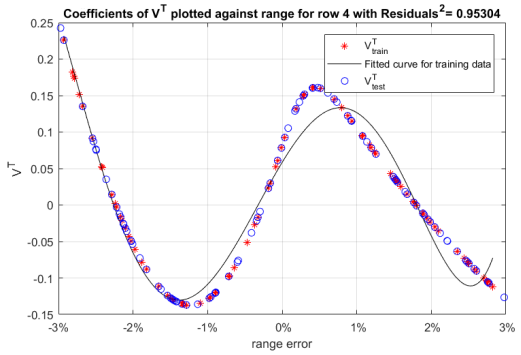

(d) Fourth row of $\boldsymbol{V}_{train}^T$ compared to $\boldsymbol{V}_{test}^T$, plotted with the regression function.

Figure 18: Comparison of coefficients from training set and test set per row.

These results on the coefficients look promising for restoring the information in the dose distribution matrix as well. The data points of the training set are on the same curve as those of the test set. The comparison for the test between the reconstructed dose distribution and the original is made in Figure 19, according to Equation (17).

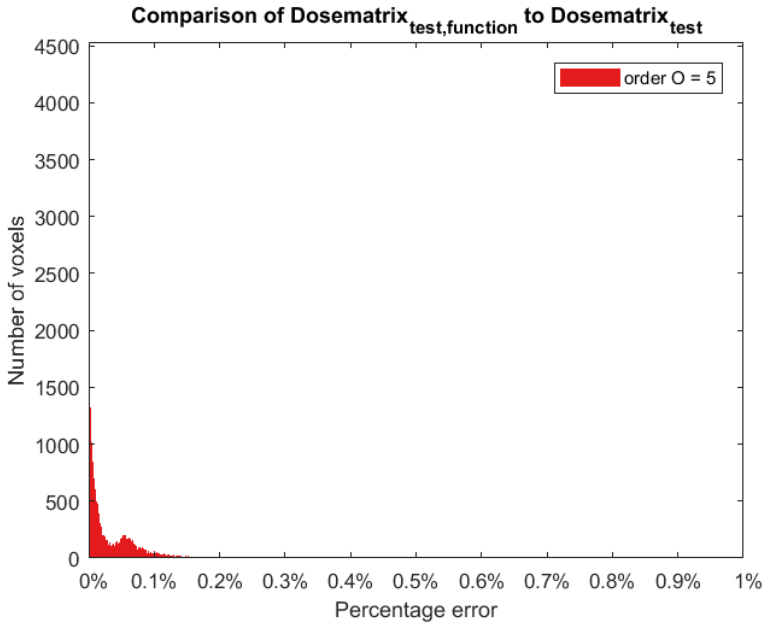

Figure 19: Comparing the test set, reconstructed with the function to the real calculated dose distribution matrix results in 91.8% acceptance of the voxels, allowing 0.1% error with voxel-by-voxel comparison. This is an average over $N = 100$ test error scenarios.

What could be expected from the coefficients of the training results occurs for the test set as well indeed. Using the reduced order and regression function works really well. The result is an almost perfectly reconstructed dose distribution matrix, accepting all of voxels, having an allowed error of 1% compared to the maximum dose. Being even stricter with an allowed margin of 0.1%, still 91.8% of the voxels are accepted, so the reconstruction results for range error scenarios are extremely good.

4.4 Positioning errors in 2 directions

After the successful procedure on range errors, position errors in 2 direction are examined. Dose masking leaves only $M_r \approx 8 * 10^4$ rows. Again 100 error scenarios are obtained.

4.4.1 Order approximation of singular values

The singular values are decreasing, just like for the range error. As can be seen in Figure 20, the values decrease with an order of magnitude of 3 within the first 11 principal components.

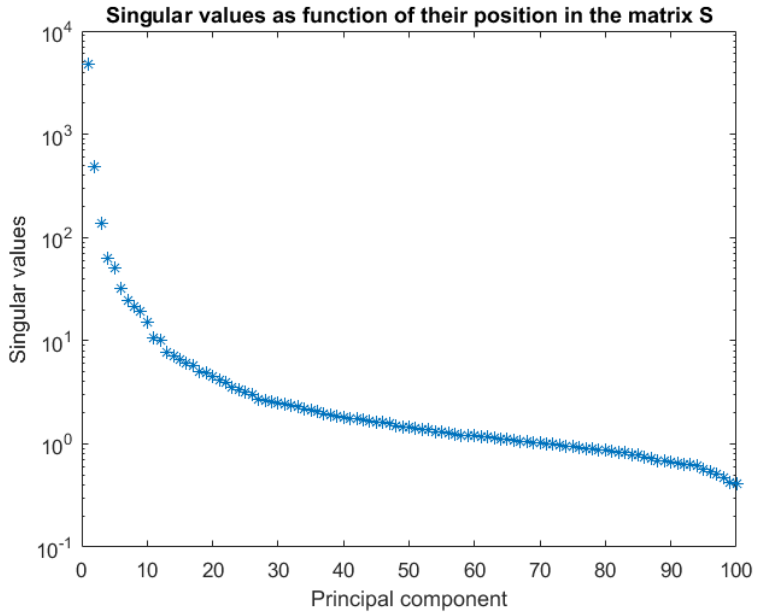

Figure 20: Singular values decrease for positioning errors in x- and y-direction, with displacements error randomly generated with a Gaussian with a standard deviation of 3 mm.

4.4.2 Reconstructing the dose distribution matrix after using a reduced order of the SVD

In Figure 21(a), the error that is made on the voxels is plotted for several orders, as an average over 100 scenarios. In Figure 21(b) this is done for the worst case scenario. From both of the images it can be obtained that taking only the first couple of orders leads to a high fraction of information being preserved.

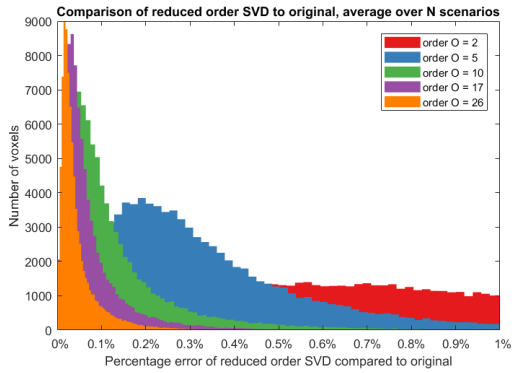

(a) A comparison between the dose distribution of an order of the SVD and the real calculated dose, as an average over the scenarios.

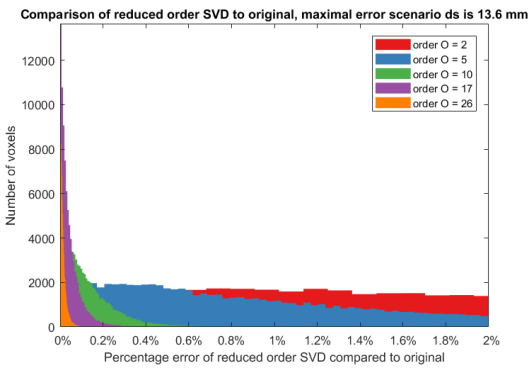

(b) A comparison between the dose distribution of an order of the SVD and the real calculated dose, for the worst case scenario, where the displacement error $ds = 13.0\,mm$.

Figure 21: Distribution of voxel dose errors for several orders, both averaged over 100 scenarios and for the worst case scenario.

In order to recompute 99% of the dose distribution matrix elements with an accuracy of 0.5%, 14 orders of the SVD are needed (Figure 22(b)). For 99.9% acceptance, the first 29 orders should be used. For a required margin of 1% in 99.9% of the voxels, an order of 12 is needed, as can be seen in Figure 22(a). In the continuation, only the first 14 principal components and first 14 eigenvectors in $\boldsymbol{V}^T$ are taken into account, to ensure accurate dose reconstruction in 99% of the voxels.

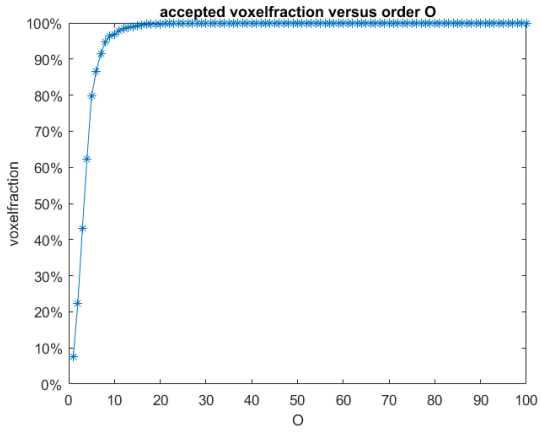

(a) Fraction of voxels having a smaller error than 0.5% per voxel, averaged over 100 scenarios, as a function of the number of orders used. $O = 14$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 29$ is required.

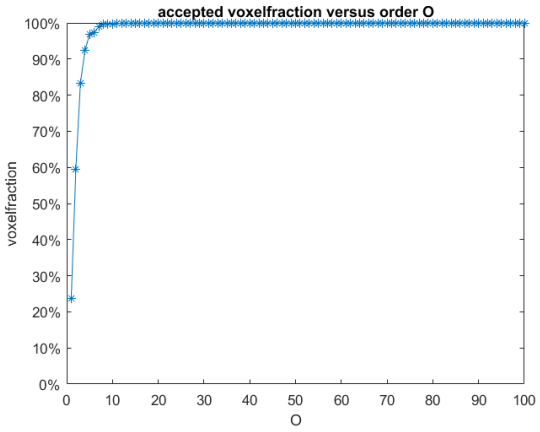

(b) Fraction of voxels having a smaller error than 1% per voxel, averaged over 100 scenarios, as a function of the number of orders used. $O = 8$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 12$ is required.

Figure 22: Accepted voxelfractions as a function of the order O, for a 0.5% (a) and a 1% (b) voxel error margin, as an average over 100 scenarios.

4.4.3 Regression

On the right singular vectors a regression is performed with polynomials up to the 5^{th} order in both x- and y-direction. The resulting surfaces can be viewed in Figure 23.

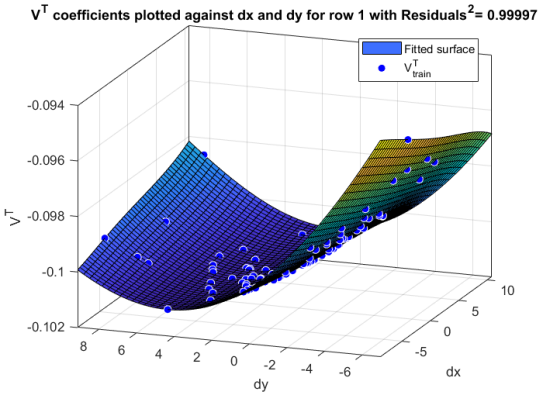

(a) Coefficients and scattered deviations with fitted surface for row 1 of $\boldsymbol{V}^T$.

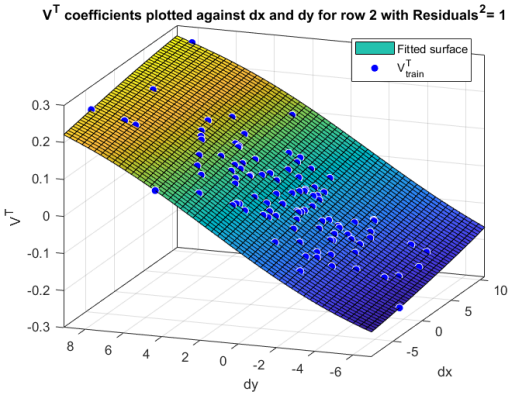

(b) Coefficients and scattered deviations with fitted surface for row 2 of $\boldsymbol{V}^T$.

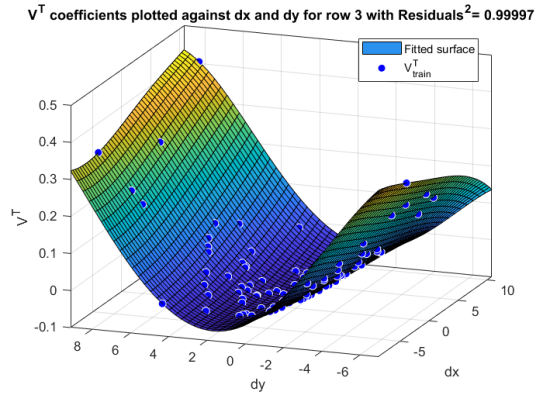

(c) Coefficients and scattered deviations with fitted surface for row 3 of $\boldsymbol{V}^T$.

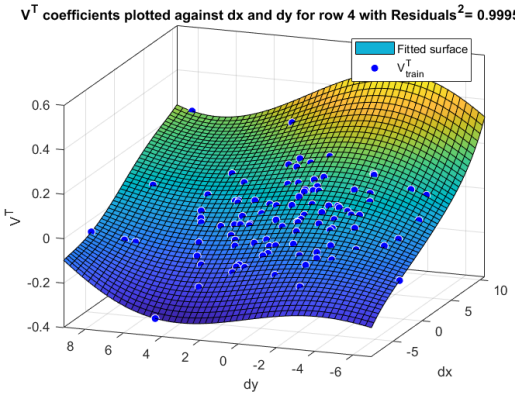

(d) Coefficients and scattered deviations with fitted surface for row 4 of $\boldsymbol{V}^T$.

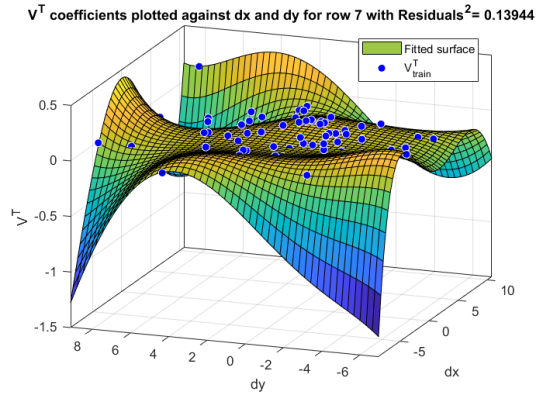

(e) Coefficients and scattered deviations with fitted surface for row 7 of $\boldsymbol{V}^T$.

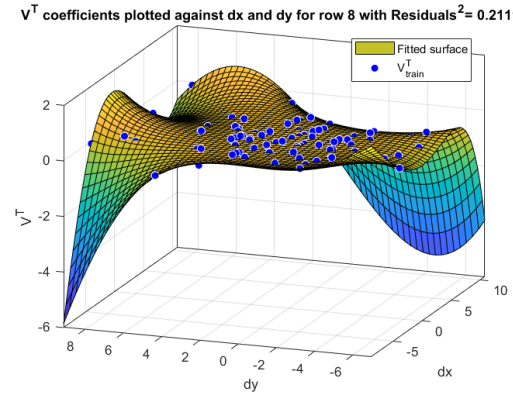

(f) Coefficients and scattered deviations with fitted surface for row 8 of $\boldsymbol{V}^T$.

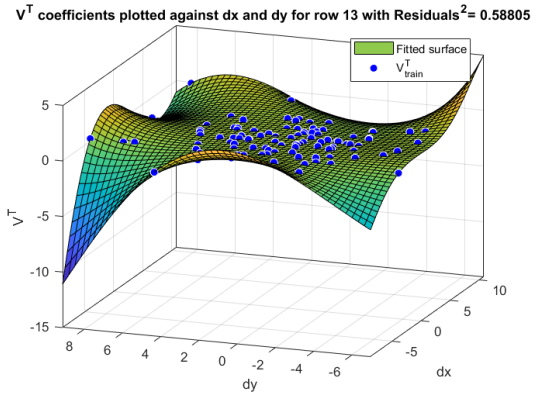

(g) Coefficients and scattered deviations with fitted surface for row 13 of $\boldsymbol{V}^T$.

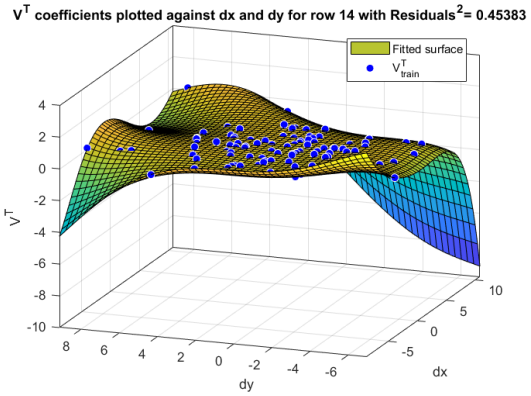

(h) Coefficients and scattered deviations with fitted surface for row 14 of $\boldsymbol{V}^T$.

Figure 23: The coefficients of $\boldsymbol{V}^T$ are plotted against the displacement data dx and dy. Also included are the fitted surfaces and corresponding residuals. In all cases, the 5^{th} order polynomials in both x- and y-direction are used.

4.4.4 Performance on train and test data

Just like in the range error case, are the resulting matrices compared. Analysing regression only is done in Figure 24(a), both the reduced model and regression in Figure 24(b).

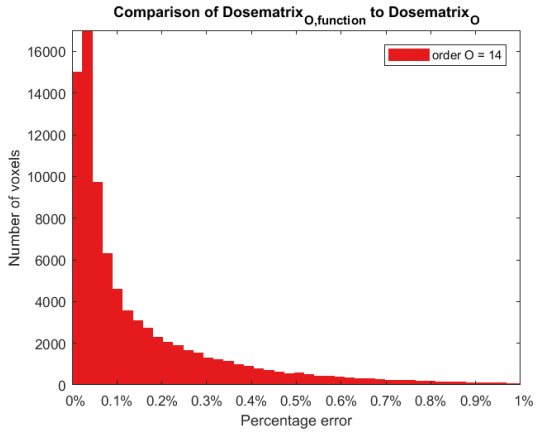

(a) Comparing the reduced order matrix with regression to the reduced order matrix results in 97.3% acceptance of the voxels, allowing 1% difference.

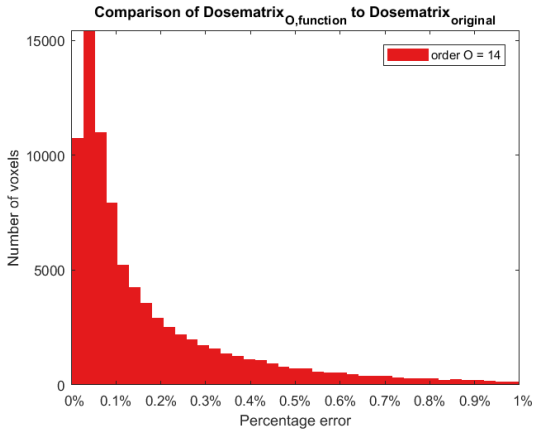

(b) Comparing the reduced order matrix with regression to the original matrix results in 96.5% acceptance of the voxels, allowing 1% difference.

Figure 24: Comparison of reduced order, regressed dose distribution matrix with the reduced order and original matrix in a histogram. The number of voxels that has a certain error is plotted for the 14^{th} order of the SVD.

The reduction of the order and subsequently regression on the right singular vectors works well on the training set, although the acceptance is lower than with range errors. A comparison is made between the coefficients for training and test data in Figure 25. It can be seen that the coefficients lay almost on the same surface, which points to a rather good working regression model.

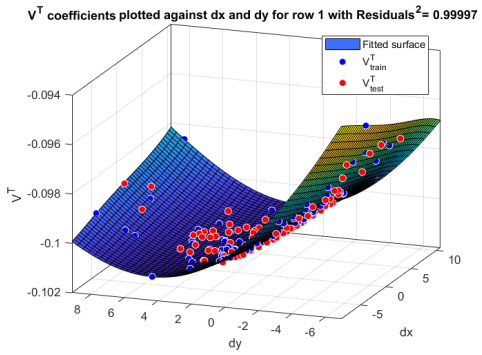

(a) Coefficients for row 1.

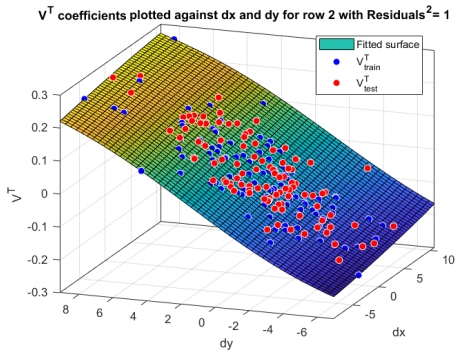

(b) Coefficients for row 2.

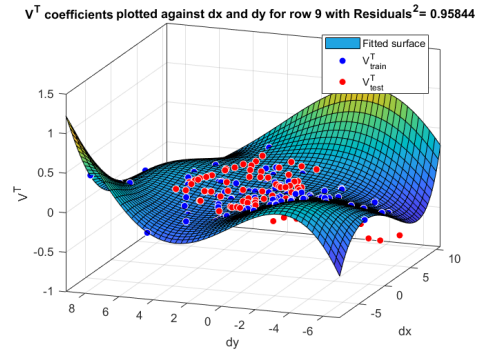

(c) Coefficients for row 9.

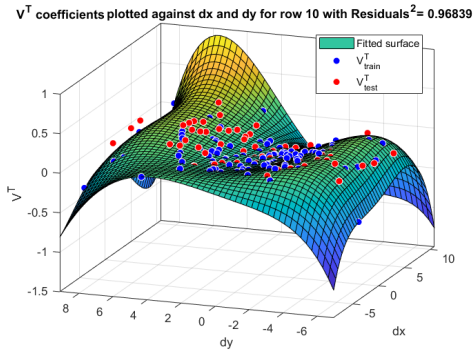

(d) Coefficients for row 10.

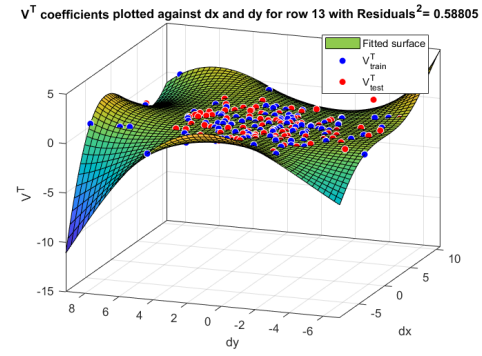

(e) Coefficients for row 13.

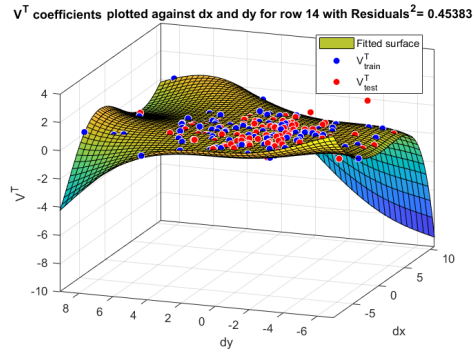

(f) Coefficients for row 14.

Figure 25: Comparison between test set and training set for several rows of $\boldsymbol{V}^T$. Also the fitted surface is plotted.

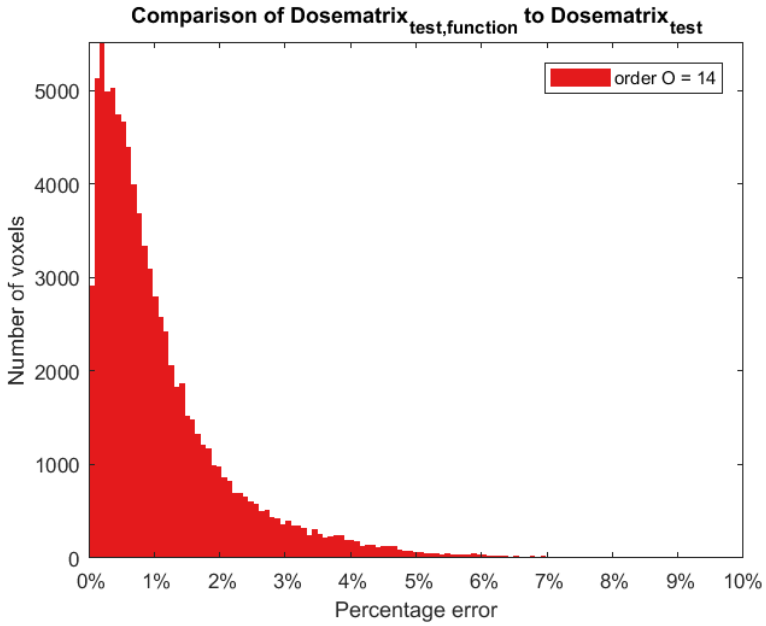

Figure 26: Comparing the test set, reconstructed with the function to the real calculated dose distribution matrix results in 64.1% acceptance of the voxels, allowing 1% error with voxel-by-voxel comparison. This is an average over $N = 100$ test error scenarios.

However, there is a significant drop in accuracy between range and positioning errors. Where for range errors all voxels were accepted, allowing 1% difference, are for positioning errors only 64.1% of the voxels within margins. It is likely that this comes from the increased number of degrees of freedom, but there might be other influences as well. In the first place, it is possible that range errors are more consistent in the influence on the dose distribution matrix. Secondly, the difference can be a result of the direction in which the displacement is performed, namely x and y. As can be seen in Figure 7, the beams travel from a gantry angle of 70° and 270°, which is for the first beam an angle of 20° to the x-axis, and for the second beam a shift parallel to this axis. A reason for the falloff can be that the regression overfits on the dependence of the error in this direction. The error in 2 direction will be repeated for directions y and z. A reason that can not be overcome by this change of direction is the influence of outliers in the test set. For that, a higher number of error scenarios would have to be simulated.

4.4.5 Positioning errors in y and z direction

For this part, some underpinning figures are included in Appendix A. The same steps that are taken for the displacement dx and dy will now be taken. In Figure 41(a), the singular values are displayed, in Figure 41 the error that is made on the voxels is plotted for several orders of the SVD. In Figure 42, the fraction of voxels that is accepted is plotted for several margins. Figure 43 shows how good the reconstruction of the original and reduced order matrix works.

To acquire same accuracy in the original dose distribution matrix as in the previous section, 21 orders have to be taken instead of 14. Subsequently, the regressions are performed and compared to Figure 23. The regression functions of Figure 44 seem to have clearer dependence on both directions, instead of a clear dependence on one direction and hardly any on the second direction. With the reduced order and regression, the original dose distribution matrix can in 98.4% of the voxels be reconstructed with an accuracy of 1%, which is low compared to the dx-dy case.

However, as can bee seen in Figure 27, the performance drops less on unseen test data, relative to training data. The regression thus works better in general case, so the problem of overfitting is reduced.

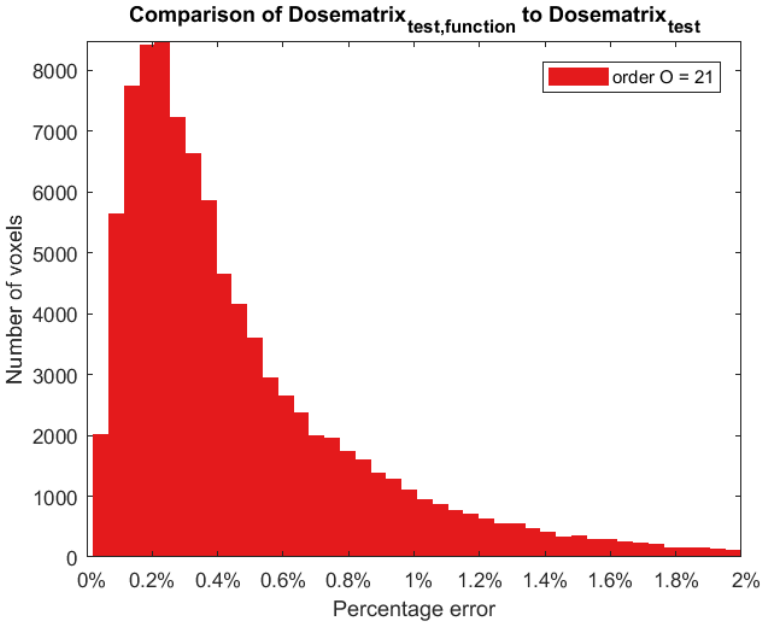

Figure 27: Comparing the test set, reconstructed with the function, to the real calculated dose distribution matrix results in 87.8% acceptance of the voxels, allowing 1% error with voxel-by-voxel comparison. This is an average over $N = 100$ test error scenarios.

4.5 Positioning errors in 3 directions

Now, heading towards an error in directions x, y and z, the matrix can be chopped from $M \approx 1.7 * 10^6$ rows to $M_r \approx 8 * 10^4$.

4.5.1 Order approximation of singular values

The decrease in singular values in Figure 28(a) is still visible. The data for these error scenarios is generated using a Gaussian with a standard deviation of $3\,mm$ in all direction. When a new data set is created with standard deviation $5\,mm$ in all direction, this drop-off tends to go a little slower, as can be seen in Figure 28(b). Still, for the first 20 orders a decrease of factor 1000 can be noticed.

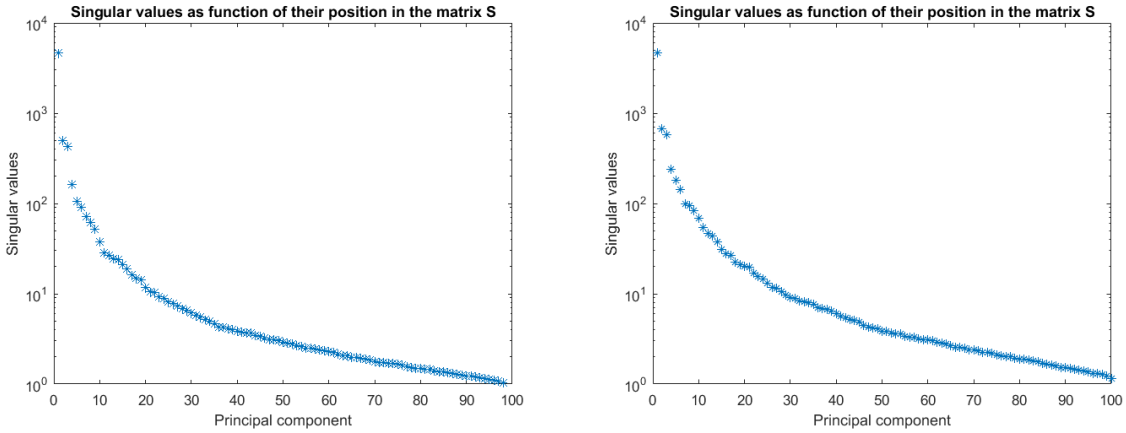

(a) Singular value decrease for the 3 directional error scenarios, where data is created with standard deviation of $3\,mm$.

(b) Singular value decrease for the 3 directional error scenarios, where data is created with standard deviation $5\,mm$.

Figure 28: Singular value decrease for two sets of 3-directional simulated errors, one with standard deviation of $3\,mm$ (a), one with standard deviation of $5\,mm$(b).

4.5.2 Reconstructing the dose distribution matrix after using a reduced order of the SVD

In Figure 29(a) and 29(b), the error that is made on the voxels is plotted for several orders, as an average over 100 scenarios. In Figure 29(c) and 29(d) this is done for the worst case scenarios. The right hand side figures represent the scenarios with bigger deviation in the displacement.

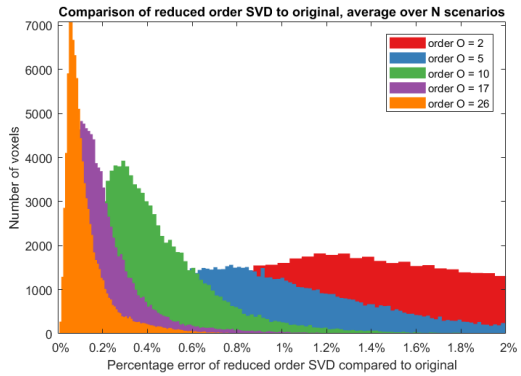

(a) Error on the dose per voxel for several orders, averaged over the 100 scenarios. Data set with positioning errors with standard deviation 3 mm.

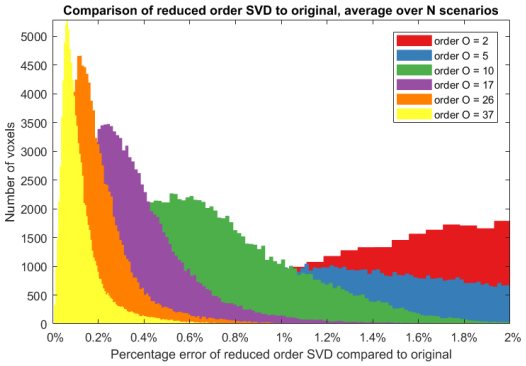

(b) Error on the dose per voxel for several orders, averaged over the 100 scenarios. Data set with positioning errors with standard deviation 5 mm.

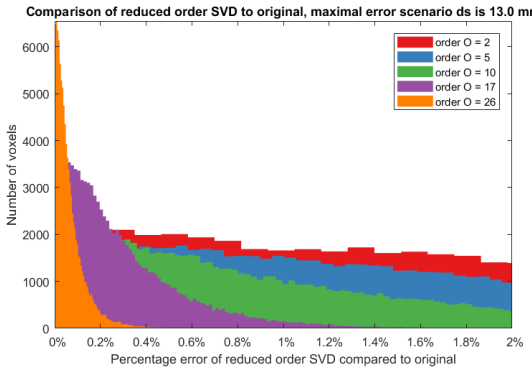

(c) Error on the dose per voxel for several orders for the worst case scenario, where $ds = 13.0\,mm$. Data set with positioning errors with standard deviation 3 mm.

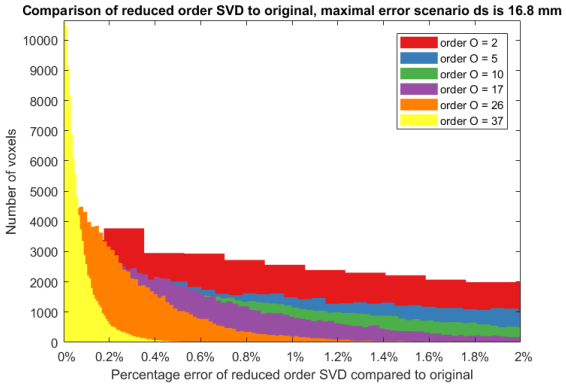

(d) Error on the dose per voxel for several orders for the worst case scenario, where $ds = 16.8\,mm$. Data set with positioning errors with standard deviation 5 mm.

Figure 29: Comparison of the influence of the order O on the accuracy of the reconstruction, both for a worst and average scenario, with positioning errors with a standard deviation of 3 and 5 mm.

In order to recompute 99% of the dose distribution matrix elements with an accuracy of 0.5%, the 31 orders of the SVD are needed, as can be seen in Figure 30(a). For 99.9% acceptance, the first 49 orders should be used. For a required accuracy of 1% in 99% of the voxels, an order of 17 is needed, as can be seen in Figure 31. In the continuation, only the first 17 principal components and first 17 eigenvectors in $\boldsymbol{V}^T$ are taken into account. For the higher-deviation case, the first 21 components and eigenvectors would be needed to reach same accuracy. A plot of the fraction of voxels that is reconstructed with an accuracy of 0.5% is shown in Figure 31.

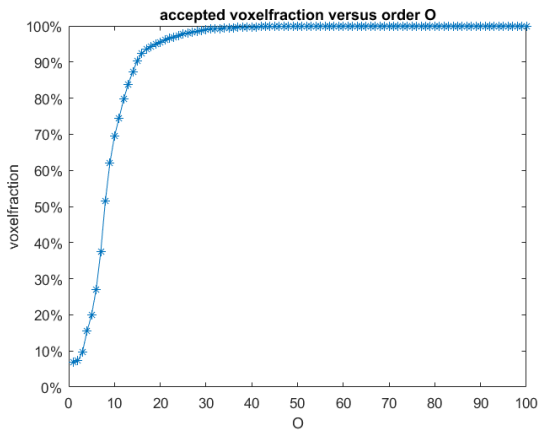

(a) Fraction of voxels having a smaller error than 0.5% per voxel, averaged over 100 scenarios with standard deviation 3 mm, as a function of the number of orders used. $O = 31$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 49$ is required.

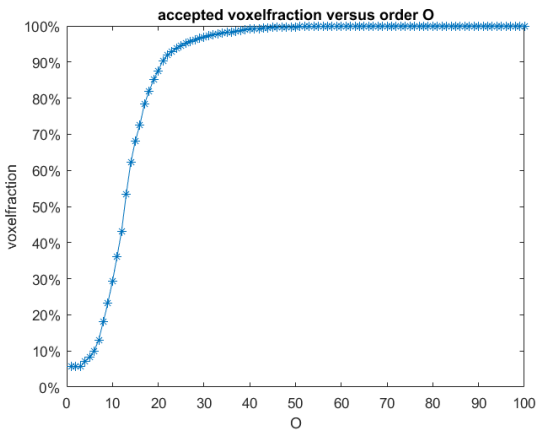

(b) Fraction of voxels having a smaller error than 0.5% per voxel, averaged over 100 scenarios with standard deviation 5 mm, as a function of the number of orders used. $O = 40$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 58$ is required.

Figure 30: The average accepted voxelfraction as a function of the order O, for datasets with a different standard deviation in the displacement error.

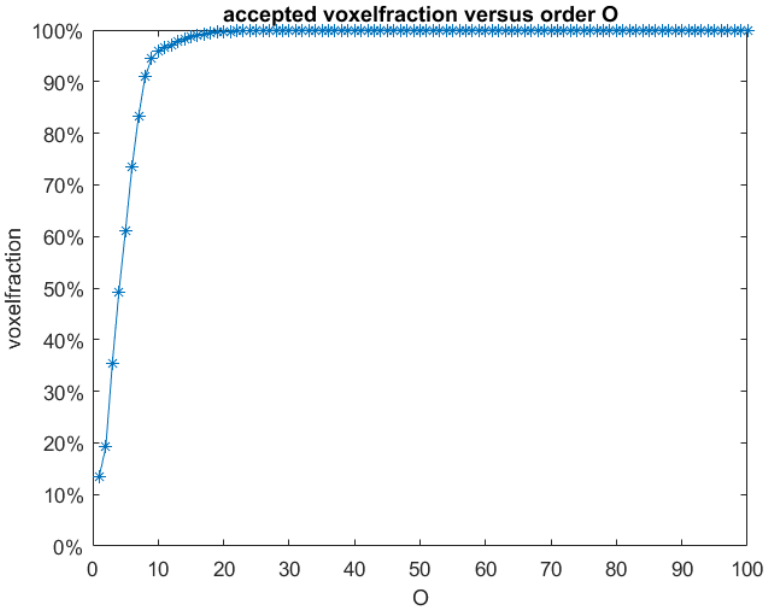

Figure 31: The average accepted voxelfraction as a function of the order O. $O = 17$ is needed when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxels, $O = 24$ is required.

4.5.3 Performance on train and test data

The result on training data are shown in Figure 32. The results are still very reasonably, quantitatively comparable to the 2 directional error scenario. Especially the accuracy on the test set, shown in Figure 33, is promising, showing 88.7% of the voxels being accepted within margins of 1%.

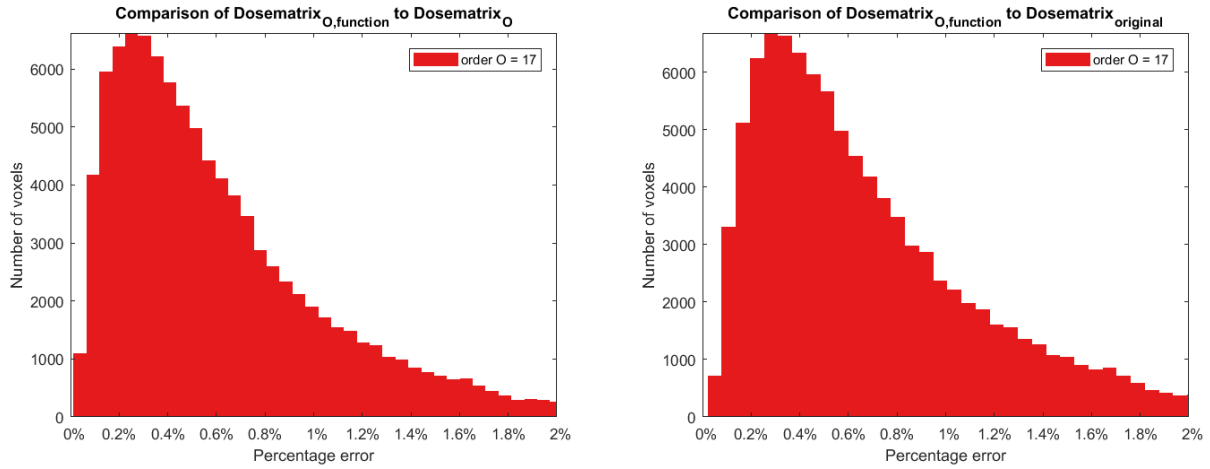

(a) Comparing the reduced order matrix with regression to the reduced order matrix results in 90.0% acceptance of the voxels, allowing 1% difference.

(b) Comparing the reduced order matrix with regression to the original matrix results in 87.1% acceptance of the voxels, allowing 1% difference.

Figure 32: Comparison of reduced order, regressed dose distribution matrix with the reduced order and original matrix.

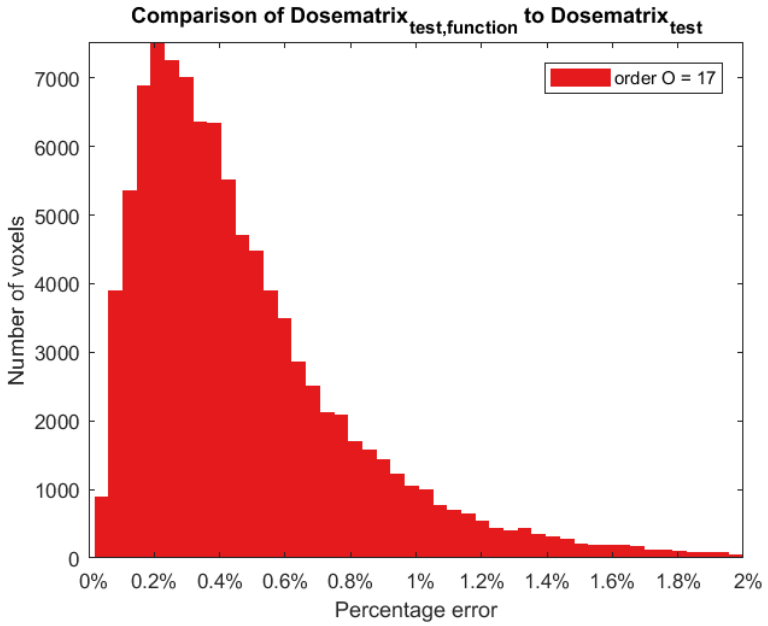

Figure 33: Comparing the test set, reconstructed with the function to the real calculated dose distribution matrix results in 88.7% acceptance of the voxels, allowing 1% error with voxel-by-voxel comparison. This is an average over $N = 100$ test error scenarios.

4.6 Combination of positioning and range errors

The more different types of errors are taken into account when simulating an error scenario, the more it represents real life scenarios. In this section, both positioning and range errors are considered. Reduction in size of the dose distribution matrix lead to a decrease from $M \approx 1.7 * 10^6$ to $M_r \approx 9 * 10^4$ rows.

4.6.1 Order approximation of singular values

The decrease of singular values, shown in Figure 34, tends to go slower than in all previous cases. However, still there is a 4 orders of magnitude difference between the first and the 25^{th} principal component.

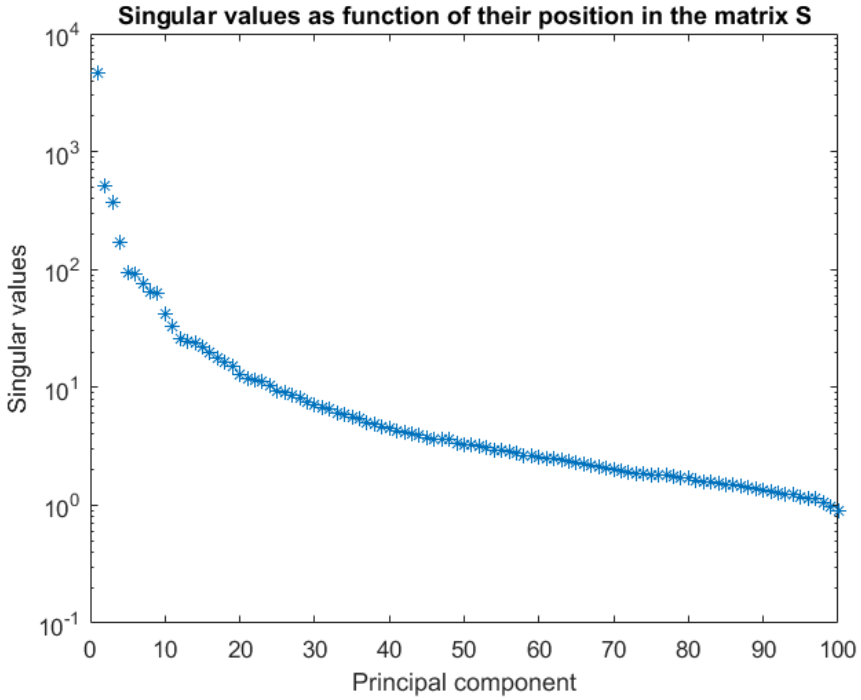

Figure 34: Singular value decrease for the scenarios with error for $d\rho$ and dx, dy and dz direction.

4.6.2 Reconstructing the dose distribution matrix after using a reduced order of the SVD

Taking too few orders will not lead to a nicely reconstructed dose distribution matrix. The dependency of the accuracy of the SVD order is shown in Figure 35 for both the average and the worst case scenarios. As can be seen, the accuracy increases when the number of orders O increases. This becomes also clear from Figure 36, from which the average accepted fraction of the voxels can be obtained.

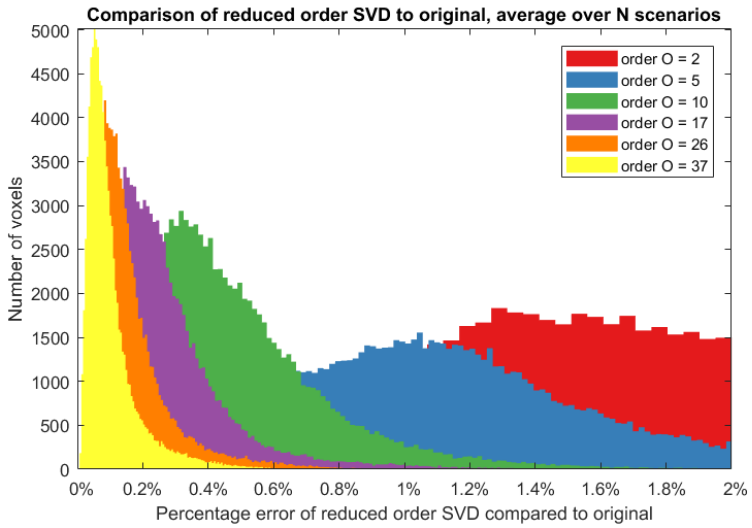

(a) Average relative error on the reconstructed voxels for several orders.

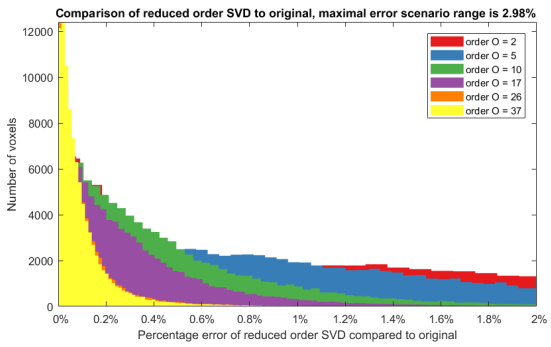

(b) Relative error on the reconstructed voxels for several orders. The error is for the worst range scenario, where $d\rho = 2.98\%$.

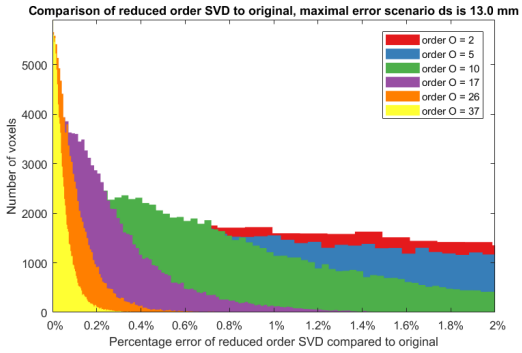

(c) Relative error on the reconstructed voxels for several orders. The error is for the worst displacement scenario, where $ds = 13.0\,mm$.

Figure 35: Histogram with the accuracy of reconstruction of the voxel doses for several orders and different scenarios.

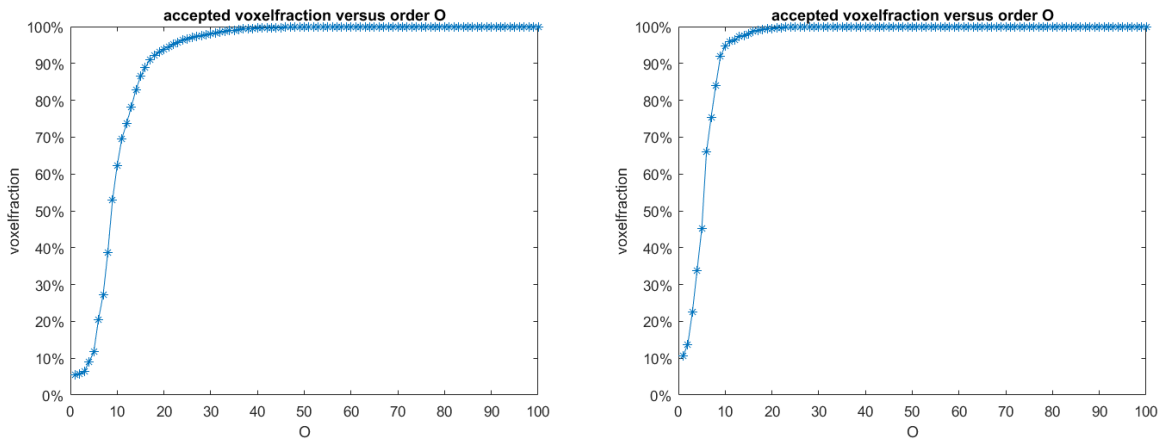

(a) 0.5% deviation per voxel requires $O = 35$ when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxel, $O = 52$ is required.

(b) 1% deviation per voxel requires $O = 18$ when 99% of the voxels should be accepted. For an acceptance of 99.9% of the voxel, $O = 26$ is required.

Figure 36: The average accepted fractions as function of the reduced order of the SVD.

4.6.3 Performance on train and test data

First, the performance of the regression alone is analysed. In Figure 37, the number of voxels that are accepted are plotted against the error. Then, the performance on the original matrix is analysed in Figure 38. As is expected, the difference between these two figures is not big, because for $O = 26$, the ROM loses only a small fraction of the dose distribution information. As can be seen, the dose distribution matrix for the training set can be reconstructed with an accuracy of 1% for 96.5% of the voxels.

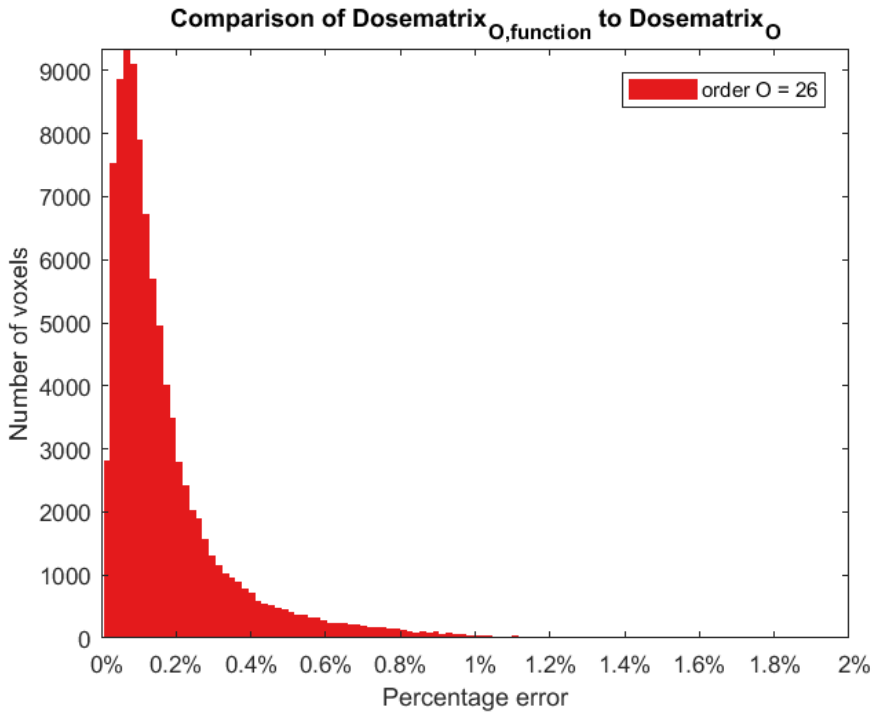

Figure 37: Comparing the reduced order matrix with regression to the reduced order matrix results in 98.0% acceptance of the voxels, allowing 1% difference.

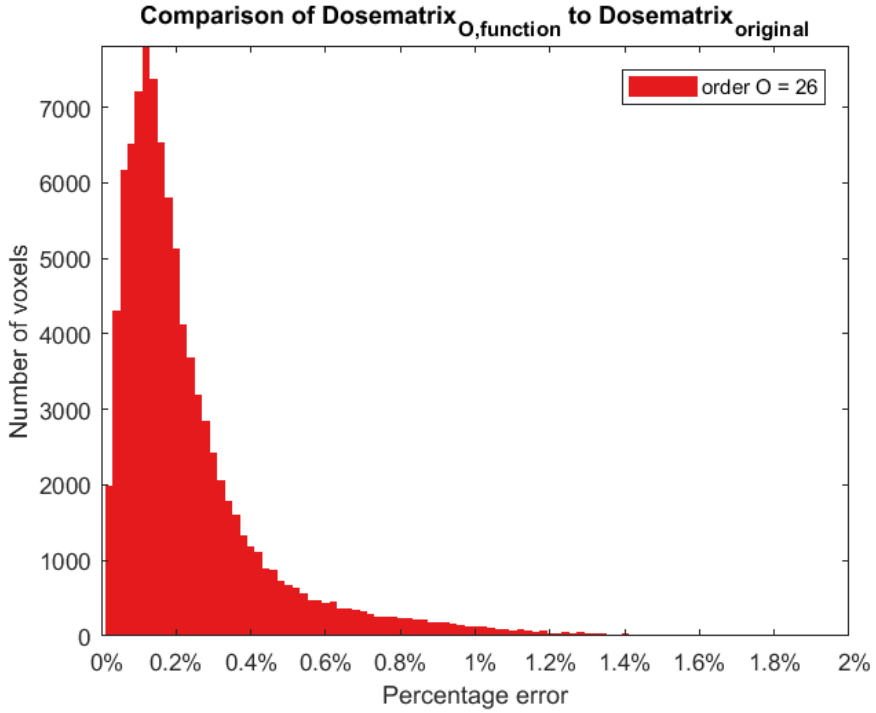

Figure 38: Comparing the reduced order matrix with regression to the original matrix results in 96.5% acceptance of the voxels, allowing 1% difference.

And again, this reduced order, regression model is tested on unseen test data. The results

show that allowing a deviation of 1%, 81.9% of the voxels are accepted, which is a reasonably high fraction. For a deviation of 3%, 92.5% of the voxels are accepted. This is a reasonable accuracy, surely regarding the computational speed of the process.

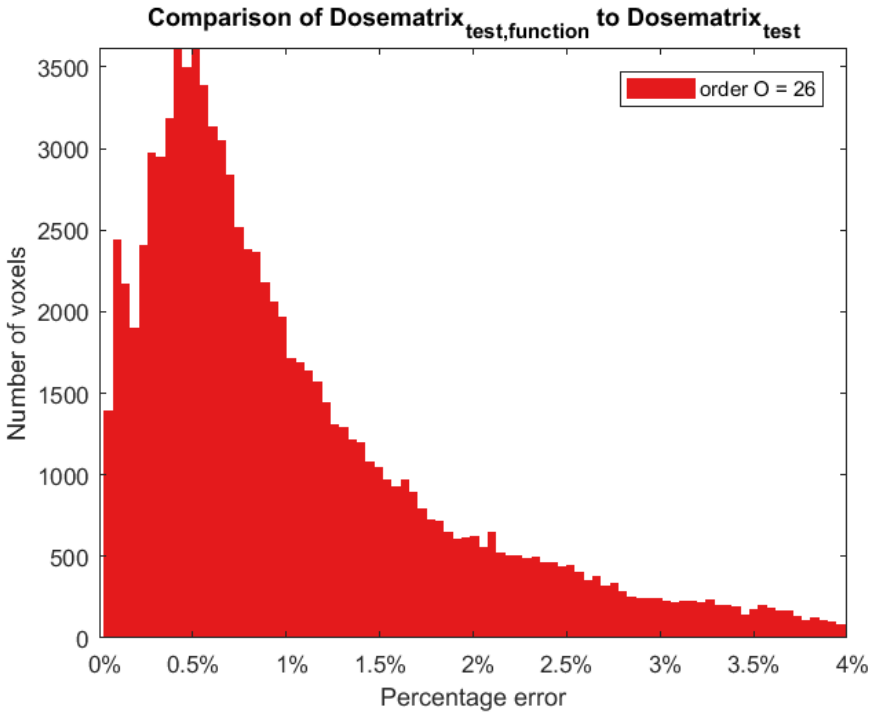

Figure 39: Comparing the test set, reconstructed with the regression function to the real calculated dose distribution matrix results in 81.0% acceptance of the voxels, allowing 1% error with voxel-by-voxel comparison. This is an average over $N = 100$ test error scenarios.

4.7 Computational speed

The results show that in most cases, the dose distribution matrix can be computed with reasonable accuracy. The time that calculating the dose distributions of 100 error scenarios took was $6 \pm 1\,h$. After this procedure was done, the reduced order model and subsequently the regression model took $0.03\,s$ for 100 error scenarios. The total time for the calculation of dose distributions for 200 error scenarios is then approximately $6\,h$. However, when the number of unseen error scenarios increases, the total computational time hardly increases. The dose distributions of any new error scenario can be calculated very quickly, using Equation (17). This is a great result for the improvement of optimisation of an IMPT plan.

5 Conclusion and Recommendations

In conclusion, it can be said that with a Reduced Order Model of the Singular Value Decomposition, a dose distribution matrix can be approximated to acceptable level. For 100 range and positioning error scenarios of at maximum 3% and a standard deviation of $3\,mm$ in all directions, only the first 35 order were needed to reach an accuracy of 0.5% margin in 99% of the voxels. Performing a regression on the right singular vectors of the singular vectors works good for range and positioning errors of at maximum 3% and a standard deviation of $3\,mm$ in all direction for the first 8 rows with a 5^{th} order polynomial function. With this reduced order, regression model, the full dose distribution matrix could be restored with 98.8% acceptance of voxel, allowing a deviation of 1%. On test data, the acceptance was lower, with 81.0% acceptance for margin of 1%. For an allowed margin of 3%, 92.5% of all voxel were accepted, using voxel-by-voxel comparison.

For future research, some topics raised interest. In the first place, the same procedure that is executed in this project could be repeated using a better nominal treatment plan, including multiple beams from multiple directions. Besides, analysing these beams could be done with a Monte-Carlo simulation instead of a pencil beam calculation.

Besides that, increasing the size of the training data set could lead to a higher accuracy in the regression, while increasing the test set lead to more certainty in results. Also, evaluating the deviations that were taken in this thesis, one can conclude that they do not represent the real treatment, where the discrepancy in the patient positioning can be reduced to no more than $5\,mm$ in each direction [16]. In order to create data that represents all of the possible error scenarios, one could also use linearly spaced intervals for the errors dx, dy, dz and $d\rho$. This reduces the influence of outliers on the regression model.

The regression on the right singular vectors is performed such that the error between $\boldsymbol{V}_f^T$ as a function and $\boldsymbol{V}^T$ is minimised. For this part, also other regression models can be examined, in order to create better fits. Alternatively, one could investigate in minimising the Euclidean norm $||\boldsymbol{U}_O * \boldsymbol{S}_O * \boldsymbol{V}_{f,O}^T - \boldsymbol{D}_{fulldose,chopped,f,O}||_2$.

In order to come up with scenarios that represent errors real treatments more accurate, different types of errors can be simulated, e.g. rotation of the patient. Also harder-to-reach error scenarios, like organ movement and tumour deformation could be taken into account.

With the whole procedure discussed in this thesis, calculating the dose distribution became computationally cheaper. However, to come up to the answer, still a number of dose influence matrices have to be calculated, before proceeding with the dose distribution calculation. Perhaps most favourable future research would be to investigate in applying the same procedure with ROM and regression on the dose influence matrix $\boldsymbol{D}_{ij}$. By doing this, the influence of the error that comes from the multiplication of the bixel weight intensities by the influence matrix can be reduced. In that case, the influence of a certain set bixel weight intensities on an error scenario could quickly be calculated with the reduced order, regressed dose influence matrix $\boldsymbol{D}_{ij,f,O}$ and the set of weights. Optimisation of the bixel weights can then be done quicker, because after each iteration, the resulting dose distribution is quickly calculated.

References

[1] Institute for Health Metrics and Evaluation. Global health data exchange. http://ghdx.healthdata.org/gbd-results-tool, 2016. Accessed: 2018-07-11.

[2] America Cancer Society. Lifetime risk of developing or dying from cancer. https://www.cancer.org/cancer/cancer-basics/lifetime-probability-of-developing-or-dying-from-cancer.html, 2018. Accessed: 2018-07-16.

[3] Martin Jermann. Particle therapy statistics in 2014. *International Journal of Particle Therapy*, 2(1):50–54, 2015.

[4] The bragg peak phenomenon. https://www.protominternational.com/2018/06/bragg-peak/, 2018. Accessed: 2018-07-11.

[5] A J Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 1: the potential effects of calculational uncertainties. *Phys. Med. Biol.*, 53, 2008.

[6] Andrew J. Wroe, David A. Bush, Reinhard W. Schulte, and Jerry D. Slater. Clinical immobilization techniques for proton therapy. *Technology in Cancer Research & Treatment*, 14(1):71–79, 2015.

[7] Harald Paganetti. Relative biological effectiveness (rbe) values for proton beam therapy. variations as a function of biological endpoint, dose, and linear energy transfer. *Physics in Medicine Biology*, 59(22), 2014.

[8] M. Arntzenius. The effect of breathing on dose uncertainties during proton therapy. 2018.

[9] H. et al. Wieser. Development of the opensource dose calculation and optimization toolkit matrad. *Phys. Med. Biol.*, 44, 2017.

[10] Linda Hong et al. A pencil beam algorithm for proton dose calculations. *Physics in Medicine Biology*, 41(8):1305, 1996.

[11] Lorenz T. Biegler Andreas Wächter. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106, 2005.

[12] Sebastian et al. van der Voort. Robustness recipes for minimax robust optimization in intensity modulated proton therapy for oropharyngeal cancer patients. *International Journal of Radiation Oncology Biology Physics*, 95, 2016.

[13] Millie Chu, Yuriy Zinchenko, Shane G Henderson, and Michael B Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Physics in Medicine Biology*, 50(23), 2005.

[14] A. C. Antoulas, D. C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. *Contemporary Mathematics*, 280:193–219, 2001.

[15] David et al. Craft. Shared data for intensity modulated radiation therapy (imrt) optimization research: the cort dataset. *GigaScience*, 3(1), 2014.

[16] Myonggeun Yoon and Cheong et al. Accuracy of an automatic patient-positioning system based on the correlation of two edge images in radiotherapy. *Journal of Digital Imaging*, 24(2), 2011.

A Figures for the dy-dz error scenarios

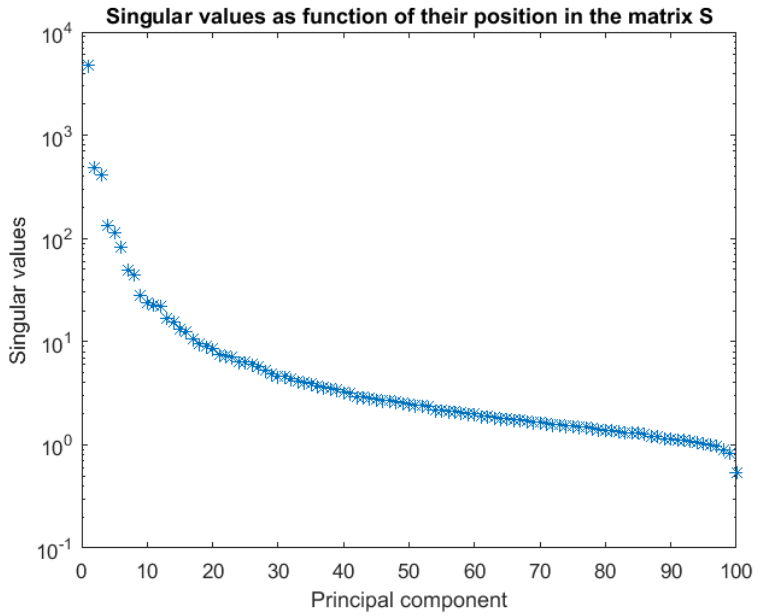

Figure 40: Singular values decrease for positioning errors in x- and y-direction, with displacements error randomly generated with a Gaussian with a standard deviation of 3 mm.

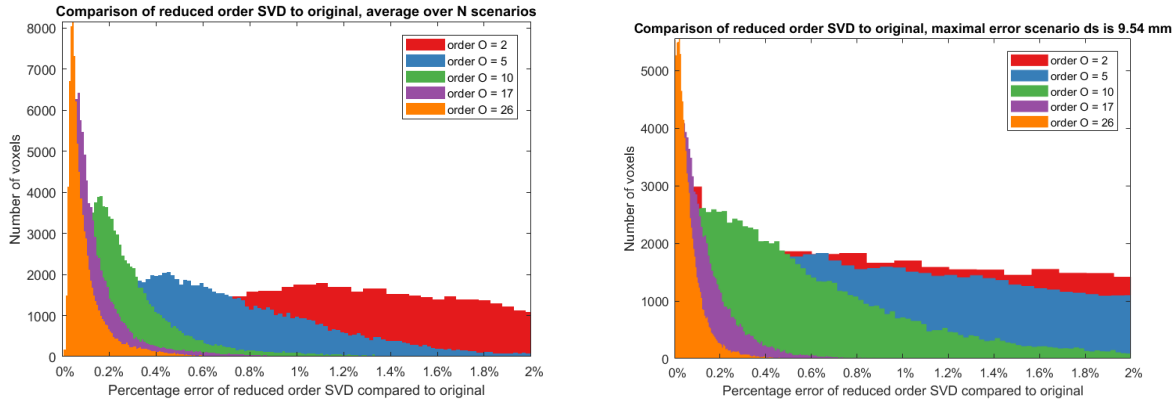

(a) A comparison between the dose distribution of an order of the SVD, compared to the real calculated dose, as an average over the scenarios.

(b) A comparison between the dose distribution of an order of the SVD, compared to the real calculated dose, for the worst case scenario, where the displacement error $ds = 9.54\,mm$.

Figure 41: Number of voxels for each accuracy plotted for several order, both for an averaged and a worst case scenario.

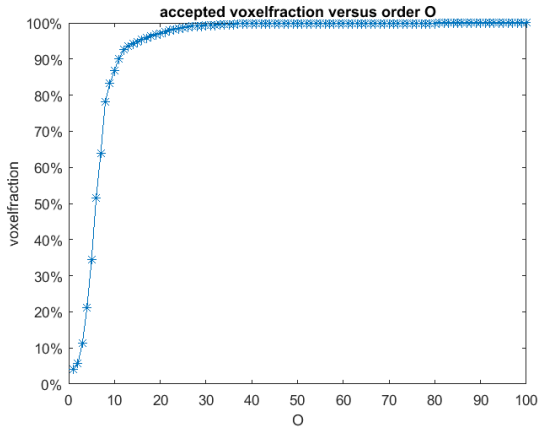

(a) Accuracy as a function of the order, with an accepted margin of 0.5% on a voxel, averaged over 100 scenarios. The 27^{th} order is needed in order to recompute 99% of the voxels within, the 44^{th} order for 99.9% acceptance.

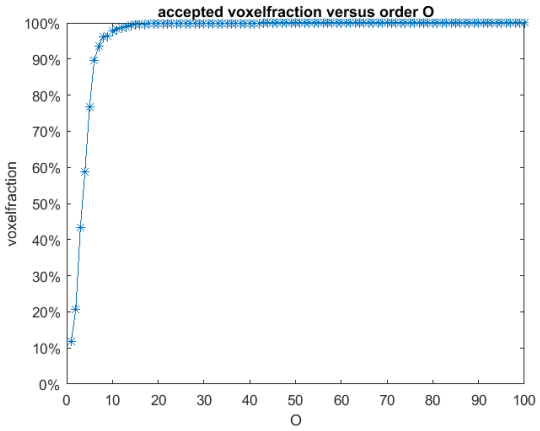

(b) Accuracy as a function of the order, with an accepted margin of 1% on a voxel, averaged over 100 scenarios. The 14^{th} order is needed in order to recompute 99% of the voxels within, the 21^{th} order for 99.9% acceptance.

Figure 42: For an average over all scenarios, the accepted voxelfraction as a function of the order, for several allowed voxel error margins.

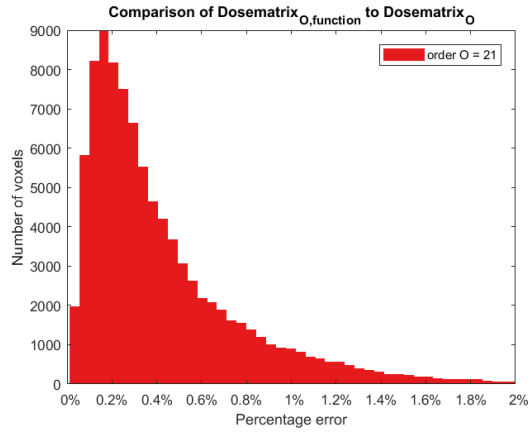

(a) Comparing the reduced order matrix with regression to the reduced order matrix results in 97.3% acceptance of the voxels, allowing 1% difference.

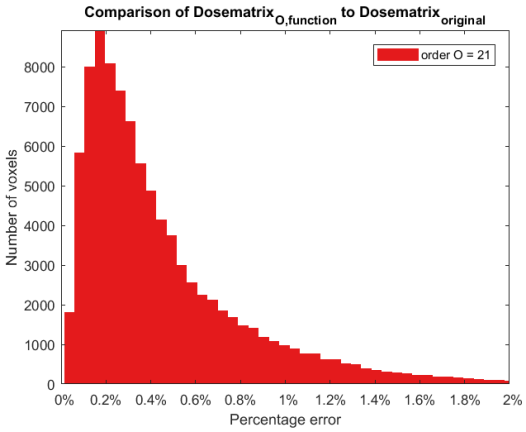

(b) Comparing the reduced order matrix with regression to the original matrix results in 96.5% acceptance of the voxels, allowing 1% difference.

Figure 43: Comparison between the reconstructed matrices and the original matrix, for the average over all scenarios. Regression is done with a 5^{th} order polynomial and $O = 21$.

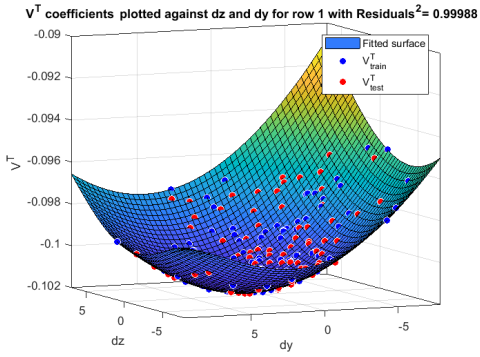

(a) Coefficients and scattered deviations with fitted surface for row 1 of $\boldsymbol{V}^T$.

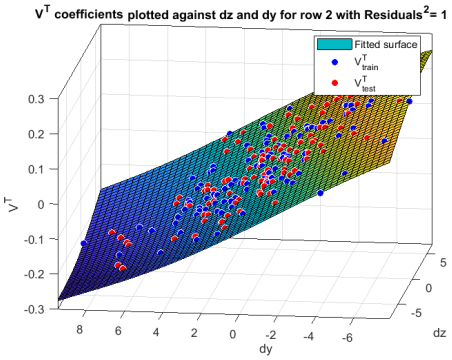

(b) Coefficients and scattered deviations with fitted surface for row 2 of $\boldsymbol{V}^T$.

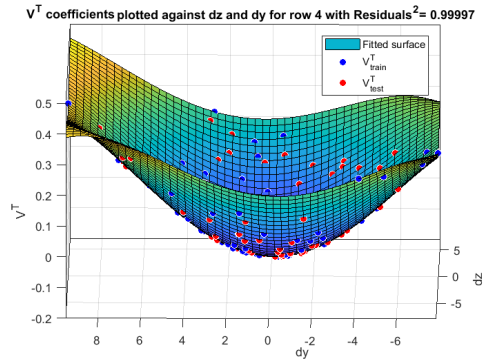

(c) Coefficients and scattered deviations with fitted surface for row 4 of $\boldsymbol{V}^T$.

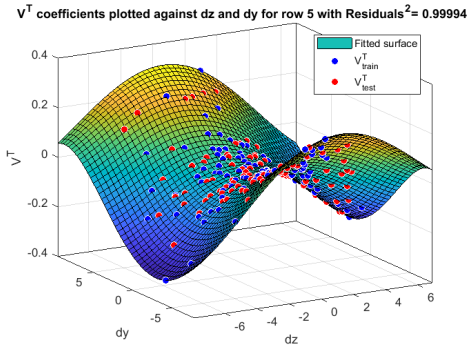

(d) Coefficients and scattered deviations with fitted surface for row 5 of $\boldsymbol{V}^T$.

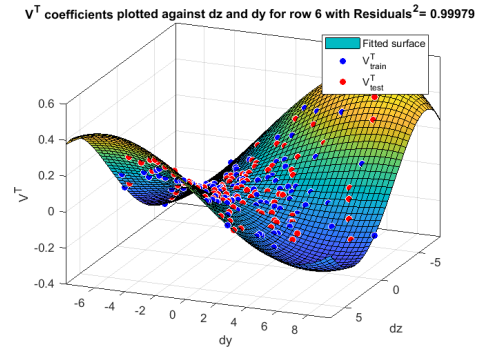

(e) Coefficients and scattered deviations with fitted surface for row 6 of $\boldsymbol{V}^T$.

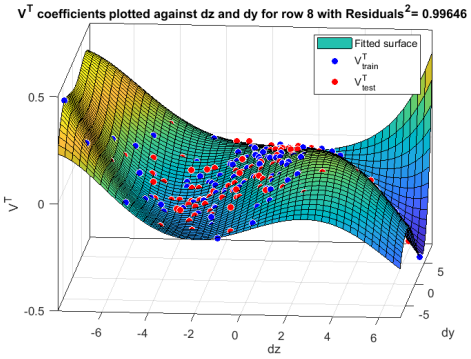

(f) Coefficients and scattered deviations with fitted surface for row 8 of $\boldsymbol{V}^T$.

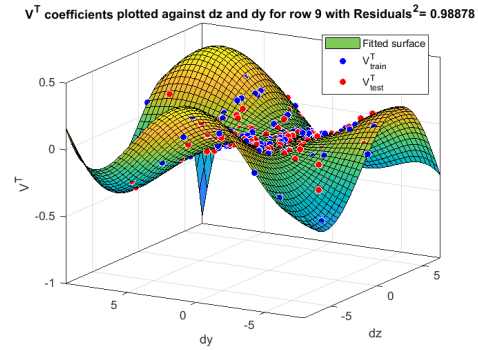

(g) Coefficients and scattered deviations with fitted surface for row 9 of $\boldsymbol{V}^T$.

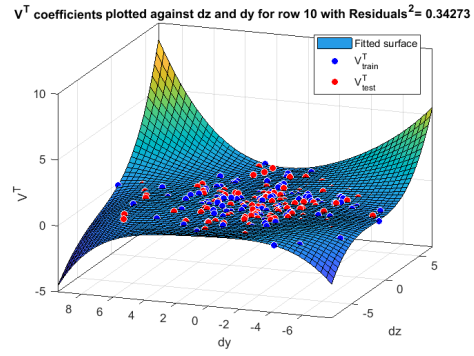

(h) Coefficients and scattered deviations with fitted surface for row 10 of $\boldsymbol{V}^T$.

Figure 44: Plots of displacement data versus the coefficients in $\boldsymbol{V}^T$, with the fitted surfaces and its residuals. In all cases, the 5^{th} polynomials is both x- and y-direction is used.