Neural network dose models for knowledge-based planning in pancreatic SBRT

Warren G. Campbell and Moyed Miften

Department of Radiation Oncology, University of Colorado School of Medicine, Aurora, CO 80045, USA

Lindsey Olsen

Department of Radiation Oncology, University of Colorado Memorial Hospital, Colorado Springs, CO 80909, USA

Priscilla Stumpf, Tracey Scheffer, Karyn A. Goodman, and Bernard L. Jones^{a)}

Department of Radiation Oncology, University of Colorado School of Medicine, Aurora, CO 80045, USA

(Received 3 May 2017; revised 3 October 2017; accepted for publication 3 October 2017; published 1 November 2017)

Purpose: Stereotactic body radiation therapy (SBRT) for pancreatic cancer requires a skillful approach to deliver ablative doses to the tumor while limiting dose to the highly sensitive duodenum, stomach, and small bowel. Here, we develop knowledge-based artificial neural network dose models (ANN-DMs) to predict dose distributions that would be approved by experienced physicians.

Methods: Arc-based SBRT treatment plans for 43 pancreatic cancer patients were planned, delivering 30–33 Gy in five fractions. Treatments were overseen by one of two physicians with individual treatment approaches, with variations in prescribed dose, target volume delineation, and primary organs at risk. Using dose distributions calculated by a commercial treatment planning system (TPS), physician-approved treatment plans were used to train ANN-DMs that could predict physician-approved dose distributions based on a set of geometric parameters (vary from voxel to voxel) and plan parameters (constant across all voxels for a given patient). Patient datasets were randomly allocated, with two-thirds used for training, and one-third used for validation. Differences between TPS and ANN-DM dose distributions were used to evaluate model performance. ANN-DM design, including neural network structure and parameter choices, was evaluated to optimize dose model performance.

Results: Remarkable improvements in ANN-DM accuracy (i.e., from > 30% to < 5% mean absolute dose error, relative to the prescribed dose) were achieved by training separate dose models for the treatment style of each physician. Increased neural network complexity (i.e., more layers, more neurons per layer) did not improve dose model accuracy. Mean dose errors were less than 5% at all distances from the PTV, and mean absolute dose errors were on the order of 5%, but no more than 10%. Dose–volume histogram errors (in cm³) demonstrated good model performance above 25 Gy, but much larger errors were seen at lower doses.

Conclusions: ANN-DM dose distributions showed excellent overall agreement with TPS dose distributions, and accuracy was substantially improved when each physician's treatment approach was taken into account by training their own dedicated models. In this manner, one could feasibly train ANN-DMs that could predict the dose distribution desired by a given physician for a given treatment site. © 2017 American Association of Physicists in Medicine [https://doi.org/10.1002/mp.12621]

Key words: artificial neural network, dose-prediction, knowledge-based planning, pancreatic cancer, stereotactic radiation therapy

1. INTRODUCTION

Pancreatic cancer is a devastating disease with an extremely high mortality. Over the past decades, treatment for early-to-mid stage pancreatic cancer has evolved significantly. The best results are seen in patients who are able to undergo surgery, with 5-yr survival rates of 20%–25% and 4%–6% with and without surgery, respectively.^{1,2} Recently, stereotactic body radiation therapy (SBRT) has emerged as a favorable option for patients with locally advanced or borderline resectable pancreatic adenocarcinoma.^{3–8} SBRT is an aggressive local therapy that has improved outcomes in other hard-to-treat tumors, such as nonsmall cell lung cancer, melanoma,

and renal cell carcinoma.^{9,10} By delivering large, ablative doses of radiation in only a few treatment fractions, SBRT leads to significantly improved local control.¹¹ However, much like surgical techniques, SBRT is a challenging form of local therapy that requires precision to achieve favorable outcomes.

Pancreatic cancer patients who undergo surgical resection of their tumor typically receive a pancreaticoduodenectomy. This procedure involves the surgical resection of the head of the pancreas, the duodenum, the gallbladder, and often the distal portion of the stomach. The remaining anatomy is then reassembled to allow bile from the liver and digestive enzymes from the residual pancreas to drain into what

remains of the small bowel, helping the patient retain some digestive function.

Our ability to safely perform this surgery today is a result of more than a century of work. Often referred to as the Whipple procedure, the origins of the modern day pancreaticoduodenectomy are often traced back to a seminal 1935 paper by Whipple, Parsons, and Mullins, in which they presented the procedure as it was performed on three patients.¹² However, the first reported pancreaticoduodenectomy for pancreatic cancer was performed nearly four decades prior in 1898 by Alessandro Codivilla for a case in which the patient died 21 days later from cachexia.¹³ In subsequent years, important developments were made before Whipple would refine the procedure, including insights into duodenal function and drainage within the gastrointestinal system.¹⁴ Still, at the time of Whipple's death in 1963, the surgery remained controversial due to its high morbidity and mortality rates.^{15,16} A 1966 review by Morris and Nardi was more hopeful, emphasizing that if mortality rates for the operation (20%–40% at the time) could be reduced, then “a more optimistic picture” of the procedure could be painted.¹⁷ Today, after decades of refinement, pancreaticoduodenectomy mortality rates have reduced substantially (as low as 1%).¹⁸ Even so, expertise and experience still play an important role. Mortality rates are significantly higher at low-volume facilities compared to high-volume facilities.¹⁹ As such, significant value must still be attributed to the skills of the surgeon.

Here, we argue that radiation therapy's role in the treatment of pancreatic cancer is following a path parallel to the path for the pancreaticoduodenectomy. Akin to broad improvements to surgical procedures in general, techniques and technologies involved in radiation delivery have advanced rapidly in recent decades. Yet, there is still significant variability between treatments at different centers. Abrams et al. recently examined the role that adherence to radiation therapy protocol played in outcomes for the Radiation Therapy Oncology Group (RTOG) 9704 phase III trial for pancreatic cancer, and they found that failure to adhere to protocol was significantly associated with reduced median survival.²⁰ In theory, inverse planning should result in treatment with an optimal plan, yet studies have shown significant variations between planners.²¹ Also, similar to surgery, Amini et al. found significant differences in survival for complex radiotherapy to the anal canal between high- and low-volume centers.²²

One avenue for refining our approach to pancreatic SBRT is knowledge-based planning, which seeks to utilize information gained from prior radiation therapy treatment plans, including the specific challenges of a given anatomical location, to help guide and improve the radiation therapy planning process. Currently, most published data regarding the efficacy of pancreatic SBRT come from single-institution studies.^{4–8} In an ideal world, it would be possible for centers with experience with this technique to disseminate data regarding how to achieve the best results. However, in practice, this is often challenging. Recently, Shiraishi and Moore developed a 3D dose-prediction model that uses artificial neural networks to predict the resulting dose distribution based on patient anatomy.²³ In this paper, we develop a

similar dose-prediction model for pancreatic SBRT, for which 3D dose-prediction is important due to the close proximity to highly sensitive organs at risk.

The purpose of this work was to test the feasibility of using knowledge-based planning for 3D dose-prediction in pancreatic SBRT. We develop artificial neural network dose models (ANN-DMs) to calculate desirable dose distributions for pancreatic cancer patients receiving arc-based SBRT. Models calculate dose according to geometric and plan parameters, and each model was trained using previous treatment plans and dose distributions at our institution. To demonstrate the appropriateness of these models, we also developed models of increased or decreased complexity. In order to guide parameter selection in future models, we also evaluated the relative importance of different parameters within the model. Ultimately, we are not advocating that the treatment approaches prescribed in this work are the best approaches for pancreatic SBRT. Rather, we intend to build a framework for the objective comparison of pancreatic SBRT plans that are meant to comply with an established set of treatment guidelines.

2. MATERIALS AND METHODS

2.A. Patient data

Data were collected from 43 patients with locally advanced, borderline resectable, or recurrent pancreatic tumors treated at our institution using arc-based SBRT. These retrospective data were collected under an internal review board-approved protocol to analyze novel methods in patient dosimetry. Planning target volumes (PTVs) were formed via a 0–5 mm patient-specific anisotropic expansion from physician-defined clinical target volumes (CTVs). Median ($\pm\sigma$) PTV was 110 (± 77) cm³. Tumors were prescribed to receive maximum doses of 30–33 Gy in five fractions, delivered using 2–4 coplanar arcs spanning 250–360 degrees. Patient immobilization and repositioning was achieved using Alpha Cradle expanding foam forming molds (Smithers Medical Products, Inc.; North Canton, OH, USA). Tumor motion was managed using either abdominal compression or respiratory gating.

Patients were treated by one of two physicians, each with an individualized treatment approach. As such, two distinct groups of patients were identified. In Group A (29 patients), two treatment volumes were specified: a larger PTV prescribed to receive 20 Gy, and a smaller, gross tumor volume (GTV) prescribed to receive 30 Gy. In Group B (14 patients), a single treatment volume, the PTV, was prescribed to receive 33 Gy. Primary organs at risk (OARs) for both Groups A and B are listed in Table I, along with contouring guidelines and dose constraints. Select OARs were contoured according to specific RTOG guidelines.²⁴

2.B. Artificial neural network dose models

Artificial neural networks can feature multiple hidden layers between an input layer and an output layer, and each

TABLE I. Organs at risk, OAR contouring guidelines, and dose constraints for patients in Group A and Group B.

Organ	Contouring guidelines	Constraint
Group A		
Stomach ^a	Includes: cardia (begins at GEJ), fundus (most cephalad, abuts left hemi-diaphragm, left and superior to cardia), body (central, largest portion), antrum (gateway to the pylorus)	Stomach minus PTV: Max point dose < 30 Gy
Bowel bag	Loops of small and large bowel and interdigitating mesentery delineated on axial CT slices. Excludes bone, muscle, separate abdominal organs (i.e., kidney, stomach, liver). Includes duodenum	Bowel minus PTV: Max point dose < 30 Gy; $V_{20} < 50 \text{ cm}^3$
Spinal cord	Contoured based on the bony confines of the spinal cord	Max point dose < 10 Gy
Liver ^a	Gallbladder should be excluded. IVC should be excluded when discrete from liver. PV should be included in liver contour when caudate lobe is seen to left of PV	Mean dose < 10 Gy
Kidneys	Both right and left kidney are contoured in their entirety	$V_{15} < 20\%$
Group B		
Stomach ^a	Same as in Group A	Stomach minus PTV: $V_{33} < 1 \text{ cm}^3$, $V_{20} < 3 \text{ cm}^3$, $V_{15} < 9 \text{ cm}^3$
Duodenum ^a	<i>First portion:</i> begins after pylorus, retroperitoneal after first ~5 cm where it is suspended by hepatoduodenal ligament. <i>Second (descending) portion:</i> starts at superior duodenal flexure, attached to head of pancreas, ~7.5 cm long, located to right of IVC at levels L1–L3. <i>Third (transverse) portion:</i> crosses in from of aorta and IVC and is posterior to SMA and SMV, ~10 cm long, marks end of C-loop of duodenum. <i>Fourth (ascending) portion:</i> travels superiorly until it is adjacent to inferior pancreatic body, ~2.5 cm long, lies anteriorly to the IMV until the IMV moves medially at the transition to the jejunum	Duodenum minus PTV: $V_{33} < 1 \text{ cm}^3$, $V_{20} < 3 \text{ cm}^3$, $V_{15} < 9 \text{ cm}^3$
Small bowel	Loops of small bowel initiating at the jejunum (end of fourth portion of the duodenum) and extending to the start of the ascending colon. Excludes the duodenum, which is contoured separately	Small bowel minus PTV: $V_{33} < 1 \text{ cm}^3$, $V_{20} < 3 \text{ cm}^3$, $V_{15} < 9 \text{ cm}^3$
Large bowel	Includes ascending, transverse, descending colon loops. Terminates at the rectum	Large bowel minus PTV: $V_{33} < 5 \text{ cm}^3$, $V_{20} < 10 \text{ cm}^3$, $V_{15} < 15 \text{ cm}^3$
Spinal cord	Same as in Group A	$V_{15} < 1 \text{ cm}^3$

TABLE I. Continued.

Organ	Contouring guidelines	Constraint
Liver ^a	Same as in Group A	Liver minus GTV: $D_{50\%} < 12 \text{ Gy}$
Kidney	Same as in Group A	$V_{15} < 35\%$

GEJ, gastroesophageal junction; IVC, inferior vena cava; PV, portal vein; SMA, superior mesenteric artery; SMV, superior mesenteric vein; IMV, inferior mesenteric vein.

^aContouring per RTOG guidelines.

hidden layer can include multiple nodes. A simple example of an artificial neural network is depicted in Fig. 1, which features two inputs, one hidden layer with three nodes, and a single output. The activation value for each hidden node is determined by taking a weighted sum of the input values (or, for multilayer networks, the set of nodes in the previous hidden layer) and then entering that weighted sum into an activation function (e.g., a sigmoid function). After activation values are calculated for each node in a layer, activation values for the subsequent layer (in this case, the output node) are calculated in a similar fashion. All weight values that interconnect nodes of one layer to nodes of the next layer are independent of one another, and these values are adjusted as the neural network is trained. Initially, all weight values are chosen randomly. Then, by providing training data (i.e., input values with known output values), error magnitudes observed at the output layer determine how weights should be adjusted, typically through backpropagation. This routine is repeated multiple times using known datasets until weight adjustments no longer have discernible effects on error size. Then, a validation dataset (i.e., input values with known output values that were not used during training) is used to assess the network's performance.

For each patient, a volume of interest was defined as all voxels within 100 mm of the surface of the PTV. Each individual voxel in the volume of interest could then be used to provide input data for an artificial neural network that would produce a single output value: dose to that voxel. Input values were either geometric parameters or plan parameters. Geometric parameters are factors such as the voxel's distance to the PTV surface, distance to an OAR, or the number of arcs directly impinging on the voxel. Plan parameters are factors such as the photon beam energy or PTV volume. Geometric parameters can differ for each voxel within a given patient, whereas plan parameters are equal for all voxels in a single patient, but can differ between patients. The initial parameters used for this work are listed in Table II and are described in more detail below. Ultimately, four main ANN-DMs were developed: a pair of models each for Group A and Group B, with a model for within the treated volume, and a model for outside the treated volume.

2.B.1. Geometric parameters

In general, all voxels were categorized into two main regions: (a) “*in slice*” voxels, which include all voxels that

FIG. 1. Left: diagram showing a simple artificial neural network with two inputs, a hidden layer with three nodes, and one output. Right: a more detailed depiction of how the activation value of a single node, a_j , is calculated. A weighted sum of all activation values from the prior layer plus a bias value is indicated here by x_j . This value is then entered into g , an activation function (e.g., we used log-sigmoid functions). Internode weight values (i.e., w_{ij} , w_{jk} , etc.) are independent of one another, and these values are adjusted as the neural network is trained.

TABLE II. Initial parameters used to build the artificial neural network dosimetric models (ANN-DMs). Geometric parameters a_{g1} through a_{g7} and both plan parameters were common to all ANN-DMs and would be generically relevant for most treatment sites receiving arc-based SBRT. Three site-specific geometric parameters were also included for patients in Groups A and B. For all parameters, minimum and maximum values are provided, along with their units.

Parameter	Name	Description	Min	Max	Units
a_{g1}	r_{ptv3D}	Shortest 3D distance to PTV	-24.7	100.9	mm
a_{g2}	r_{ptv2D}	Axial component of r_{ptv3D}	-25.0	100.0	mm
a_{g3}	z_{in}	Relative Sup-Inf position, in slice	0	1	Normalized
a_{g4}	z_{out}	Sup-Inf position, out of slice	0	1	Normalized
a_{g5}	r_{surf}	Depth from patient surface	0	109.2	mm
a_{g6}	F_{arc}	Arc factor	0	1	Normalized
a_{g7}	θ	Angle relative to PTV	-3.14	3.13	Radians
a_{p1}	V_{ptv}	Target volume	19.67	175.50	cm ³
a_{p2}	E	Photon Energy	6	10	MV
Group A					
a_{g8}	r_{st}	Shortest distance to stomach	-23.6	165.5	mm
a_{g9}	r_{bb}	Shortest distance to bowel bag	-34.1	105.5	mm
a_{g10}	r_{gtv}	Shortest distance to GTV	-13.7	110.5	mm
Group B					
a_{g8}	r_{st}	Shortest distance to stomach	-15.9	130.2	mm
a_{g9}	r_{duo}	Shortest distance to duodenum	-12.1	121.5	mm
a_{g10}	r_{sb}	Shortest distance to small bowel	-26.1	130.0	mm

lie in an axial slice that contains at least some portion of the PTV, and (b) “out of slice” voxels, which include all other voxels.

The primary parameter in the geometric model is r_{ptv3D} , the shortest 3D distance from the voxel to the PTV surface. This factor was meant to capture the general shape of the dose gradient outside the target. A related parameter, r_{ptv2D} , describes the axial distance from the voxel to the PTV surface (i.e., the shortest distance ignoring the superior–inferior displacement). Inclusion of this factor allows for more accurate differentiation in the model between voxels inside and outside treated slices (i.e., slices that include at least a portion of the treatment volume). Parameter z_{in} is the normalized superior–inferior (SI) distance from the voxel to the PTV centroid for voxels within the treated slice (i.e., ranging from 0 for voxels in the central PTV slice to 1 for voxels in the most superior or most inferior PTV slice). This factor captures photon scatter effects between slices. Parameter z_{out} is the normalized SI distance to the PTV (normalized relative to PTV height) for voxels out-of-slice with the target, and this helps to further characterize the scatter distribution outside of the direct beam. The depth parameter, r_{surf} , is the distance from the voxel to the patient surface along the axis between the voxel and the PTV. The arc factor, F_{arc} , is a binary parameter that describes whether a voxel lies in the direct path of an incoming treatment beam. Most plans examined in this study used 360° arcs, but some plans used anterior arcs that do not directly irradiate posterior regions such as the kidneys or spinal cord. The parameter θ is the angle within the axial plane, calculated with respect to the PTV centroid.

Three geometric parameters specific to the pancreas were also included in each patient group for better dose model performance. The majority of these parameters were selected based on the OARs nearest the target volumes. Both Group A and Group B included r_{st} , the shortest distance between the voxel and the stomach volume. Group A also included a second OAR factor, r_{bb} , which is the shortest distance between the voxel and the “bowel bag” volume. An additional factor in Group A, r_{gtv} , describes the shortest distance between the voxel and the GTV, and captures the effect of the multiple dose prescription levels in Group A plans. Group B patients included two additional OAR factors: r_{duo} and r_{sb} , which respectively are the shortest distances between the voxel and

the duodenum volume and small bowel volume. For all distance-to-volume-based geometric parameters, distance values are negative when voxels occur within the volume in question.

2.B.2. Plan parameters

Two plan parameters were chosen to capture broad differences between plans for different patients. The first parameter, V_{ptv} , denotes the volume of the PTV in cm^3 . The second parameter, E , indicates the photon energy used for treatment, either 6 MV, 10 MV, or 10 MV in flattening filter free (FFF) mode. Two additional plan parameters were eliminated when patients were divided into two separate groups. First eliminated was R_x , which indicated the prescribed maximum dose. Second eliminated was MD , which indicated the identity of the approving physician. The factor MD was initially considered in order to account for differences in physician treatment style, including the types of hot spots considered permissible, OAR dose-volume histogram priorities, and relative tradeoffs between PTV coverage and OAR sparing. Separating patients into Groups A and B effectively incorporates these two factors into the dose models without having to explicitly indicate them for each patient.

2.B.3. Model training and validation

Multiple ANN-DMs were developed, but all were trained and validated using data exported from the TPS (Eclipse; Varian Medical Systems; Palo Alto, CA, USA), including CT images, treatment plans, structures, and TPS-calculated dose distributions. For each individual voxel, inputs to the ANN-DM were geometric and plan parameters (as described above), and the single output was a prediction of the TPS-calculated dose for that voxel. By inputting thousands of voxels across multiple patients, prediction models are established for each ANN-DM to estimate physician-approved dose.

Patient datasets were randomly divided into two groups: roughly two-thirds for training, and one-third for validation (19 and 10 for Group A, 9 and 5 for Group B, respectively). Only voxels within 100 mm of the PTV were used as inputs. Across all 43 patients, a total of 75,610,117 voxels were available for training and validation, with 8,12,908 of those voxels residing within the PTV. However, due to the highly correlated nature of TPS dose distributions, the number of truly unique data points is likely smaller. To lower computation times, a random subset of 1 in every 40 voxels (2.5% of the total) was used for training. This choice was validated by also training models with 5%, 10%, and 25% of the total voxels, and no significant changes were seen in the results. For each group of patients, two ANN-DMs were computed: one for voxels inside the PTV, and another for voxels outside the PTV. All computations were performed in the MATLAB computing environment (MathWorks; Natick, MA, USA). To produce a model that was most accurate in the area close to the PTV, voxels were weighted by dose (i.e., $weight = dose/prescription$) to ensure sufficient sampling of the dose falloff region.

The majority of ANN-DMs used in this study were composed of a feed-forward network with 25 nodes in a single hidden layer, and they were trained using L_2 regularization. Log-sigmoid functions were used for all activation functions, and scaled conjugate gradient backprojection was used to train each network. Other ANN training algorithms were investigated (i.e., quasi-Newton backprojection and conjugate gradient backprojection), but no significant improvements to the results were seen. To evaluate the role of neural network complexity, multiple networks with increased and decreased complexity were also tested using a consistent set of model parameters. Neural network complexity increases with increasing number of hidden layers and increasing number of nodes per layer. Here, we tested neural networks with 1–3 hidden layers, with each layer containing 10 to 50 nodes in each layer. Mean absolute dose error was used to quantify any benefits gained with increasing complexity.

2.B.4. Gaussian broadening of target volumes

In early efforts to train the models, we observed good performance within the training dataset, but very poor performance in the validation dataset. These errors were found to originate mainly from overfitting of the volume parameter, V_{ptv} , which is a continuous variable, but takes a discrete value for all voxels belonging to a single patient. To correct for this, Gaussian noise was added to the V_{ptv} value seen in each voxel for a given patient. The standard deviation of this distribution was chosen empirically to be 10 cm^3 , as this was found to generate nearly continuous coverage of the volume parameter space across all patients. In this way, we were able to prevent strange model behavior in the validation dataset.

2.C. Evaluation of model performance

To quantify the validity of each dose model, metrics were calculated by comparing dose distributions calculated by TPS and dose distributions determined by ANN-DMs. For each voxel in the dose distribution, the dose difference $D_{\text{model}} - D_{\text{tps}} = \Delta D$ was calculated. Using the shortest 3D distance to the PTV surface, $r_{\text{ptv}3D}$, voxels were binned in 2 mm increments to differentiate errors in different regions of the dose gradient. Standard descriptive statistics were calculated to examine differences in the distribution of errors in training and validation datasets. Additionally, modeled dose distributions were used to calculate dose-volume histograms (DVHs) for each OAR in the model. DVH error was defined as the difference (in cm^3) between the modeled DVH and the TPS-determined DVH at each point in the DVH curve.

2.D. Simple dose model comparisons

For comparison, two additional, simpler dose models were also tested using the same data to evaluate the performance of each ANN-DM. The first and most simple of the two models was the null model, which assumes that all voxels within a volume receive the prescribed dose to that volume (e.g.,

33 Gy within the PTV, and 0 Gy outside the PTV). The second model uses simple linear regression with feature engineering, using the same set of input parameters used in the final versions of each ANN-DM.

3. RESULTS

It became evident early on in ANN-DM development that separate models needed to be created for patients from Group A and Group B in order to account for the distinct treatment approaches taken by each group's respective physician. To demonstrate this, Fig. 2 illustrates the substantial improvements in accuracy that were gained by creating dedicated models for each group. Mean absolute dose errors of $> 30\%$ in the dose falloff region were reduced to $< 5\%$ when each group was modeled separately. After preliminary investigations, several different dose model configurations were implemented in order to evaluate choices of ANN-DM structure. Two main aspects were considered: neural network complexity and choices of input parameters.

3.A. Model complexity

Mean absolute dose error with respect to model complexity was evaluated by varying the number of hidden layers in the network (1–3) and the number of nodes per layer (10, 20, 30, 40, and 50). Increased neural network complexity (i.e., more hidden layers and nodes per layer) did not significantly improve model performance. Although some reduction in errors could be seen in the training dataset as the number of hidden layers increased, those reductions were not reflected in the validation dataset. Ultimately, a consistent model structure was chosen to include one hidden layers with 25 nodes

FIG. 2. Substantial reductions in ANN-DM dose error were achieved by creating dedicated models for each distinct group of patients. Box plots (mean, 95% CI) of dose errors in validation datasets are shown with respect to distance from the PTV for All Patients, Group A, and Group B. Most notable in the dose falloff region, mean absolute dose errors were decreased from $> 30\%$ to $< 5\%$ when groups were modeled separately. [Color figure can be viewed at wileyonlinelibrary.com]

each. Using this structure, each ANN-DM took roughly 5 min to train on a standard Intel Core-i7 desktop, or roughly 20 s on a commercial-grade GPU (Nvidia Quadro M6000).

3.B. Relative pertinence of model parameters

To quantify the relative impact each parameter had for each dose model, parameters were individually removed from each model to see how much its absence increased that model's mean absolute dose error. Dose errors were summarized for three different regions: inside of the PTV, within 3 cm just outside of the PTV, and more than 3 cm outside of the PTV. Results from these tests are provided in Table III. The most impactful parameter was z_{out} , which saw a 40% increase to mean absolute dose error in the region just outside of the PTV. Nevertheless, all other parameters included in Table III showed considerable pertinence in at least one of the three regions for either Group A or Group B.

Two input parameters from Table II were not included in Table III because it was decided to remove them from all dose models: E and θ . Removal was either due to the parameter being redundant, or its removal from dose models resulted in improved accuracy in all three regions. For Group A, removal of the E factor reduced the model's overall error by roughly 4.5% of the prescribed dose. For Group B, all treatment plans used 10 MV FFF beams, so inclusion of the E parameter was unnecessary. For both groups, inclusion of θ also resulted in increased errors in all regions.

3.C. Model accuracy

Using the final versions of each dose model (i.e., with E and θ removed), sample data are shown in Fig. 3 for two patients, one from Group A and one from Group B. Included

TABLE III. Quantifying the relative contribution of each parameter to model accuracy. Values shown are the relative increase in mean absolute dose error (as % of Rx dose) when each parameter was removed from the dose model. Analysis was divided into three regions: inside the PTV, within 3 cm outside of the PTV, and more than 3 cm outside of the PTV.

Parameter	Inside PTV		≤ 3 cm Outside PTV		> 3 cm Outside PTV	
	A	B	A	B	A	B
r_{ptv3D}	1.8	11.8	4.9	12.8	0.9	15.6
r_{ptv2D}	3.8	5.3	13.5	14.0	5.0	0.9
z_{in}	2.8	4.8	12.5	8.4	5.8	4.9
z_{out}	1.6	24.2	40.2	37.6	10.5	2.8
r_{surf}	7.0	7.7	1.7	6.9	-0.8	12.4
F_{arc}	0.1	5.7	8.6	1.6	3.8	-7.2
V_{ptv}	0.0	11.3	6.7	3.3	1.4	-5.2
r_{st}	2.8	4.6	1.1	8.6	-1.4	9.7
r_{bb}	0.7	—	9.9	—	10.3	—
r_{gtv}	31.2	—	2.5	—	1.2	—
r_{duo}	—	8.7	—	0.2	—	-4.2
r_{sb}	—	-1.2	—	10.8	—	1.3

FIG. 3. Sample patient data using final versions of each dose model. From Group A: (a) an axial slice from the planning CT, (b) the corresponding dose error map, and (c) the dose error map displayed with respect to angle and distance to the PTV centroid. The same data from Group B are respectively shown in (d), (e), and (f). [Color figure can be viewed at wileyonlinelibrary.com]

are axial slices from the planning CT, their corresponding axial dose error maps, as well as dose error maps displayed with respect to angle and distance to the PTV centroid. Across all patients, the maximum and minimum dose outputs for our ANN-DMs were 36.38 Gy and -1.25 Gy, respectively.

Three performance metrics — mean dose error, mean absolute dose error, and dose error standard deviation — were calculated for training and validation data for Group A and Group B, and are plotted in Fig. 4 with respect to voxel distance to the PTV surface. In order to present the entire dose fall off region, values just within the PTV surface (i.e., negative distance from PTV values) are also shown. Overall, modeled dose distributions displayed excellent agreement with planned dose distributions. Mean dose errors were less than 5% at all distances from the PTV, while mean absolute dose error was on the order of 5%, but no more than 10%. The magnitude of errors was similar between Group A and

Group B, although the mean error and mean absolute error were both higher inside the PTV in Group A, possibly due to the multiple dose levels used. The standard deviation of the dose distribution quantified the distribution of errors, and in both groups was similar in the training and validation datasets.

DVH errors for training and validation data are shown in Fig. 5 for each of the primary OAR volumes used for Group A and Group B. Models displayed good performance in higher dose regions of the DVH curve, but much larger errors were seen at lower doses. Based on the OAR dose restraints used in this work (see Table I), only the V_{33} metric could be reliably predicted.

3.D. Simple dose model comparisons

Voxel error distributions are plotted in Fig. 6 for the final versions of each ANN-DM, the null model, and simple linear

FIG. 4. Dose errors from final versions of each dose model. Mean dose errors, mean absolute dose errors, and dose error standard deviations are shown separately for Groups A and B, plotted as a function of distance from the PTV. Voxels just within the boundaries of the PTV (indicated by negative distances from the PTV) are also included to sample the full range of the dose fall off region. Box plots indicate mean values and 95% CI ranges. [Color figure can be viewed at wileyonlinelibrary.com]

regression models. For each, training and validation data are both shown. Outside the PTV, dose distributions based on linear regression dose models demonstrated much broader error ranges than ANN-DM dose distributions. Within the PTV, the null model was shown to underestimate the dose delivered, indicating that a noticeable region of the PTV receives more than the nominal prescribed dose in physician-approved clinical treatment plans.

4. DISCUSSION

We have implemented a neural-network-based 3D dose-prediction algorithm, and have trained/validated it using 43 clinically approved pancreatic SBRT plans. The algorithm is

able to predict 3D dose distributions in the vicinity of pancreatic tumors with errors on the order of 5%. ANN-DMs were also shown to significantly outperform more simple models (i.e., null model and simple linear regression). However, there remains substantial room for improvement of the models. One advantage of neural network machine learning is that new information can be incorporated with relative ease by the additional of parameters to the model. We note that, in Figs. 3c and 3f, there is a substantial component of error related to angle, although the addition of a single angle (i.e., θ) parameter was insufficient to account for this difference. One possibility is that the model may require additional information regarding the path of each radiation beamlet (e.g., whether there are overlying PTV or OAR voxels).

FIG. 5. Each patient's model-calculated dose distribution was analyzed in its entirety to calculate dose-volume histograms (DVHs) for primary organs at risk. Here, DVH errors are shown (in cm^3) for each primary organ-at-risk in Groups A and B. [Color figure can be viewed at wileyonlinelibrary.com]

Our results demonstrate the appropriateness of artificial neural networks for this purpose. Compared to more simple models (null models or linear regression with feature engineering), ANN-DM displayed greatly reduced error (see Fig. 6). The underlying structure of planned radiotherapy dose distributions is inherently multidimensional, and modeling these distributions requires a sophisticated approach. Additionally, we found no benefit to more complex models (e.g., increasing the complexity of the neural network structure).

As has been demonstrated, artificial neural network dose models can be used to reliably predict physician-approved dose distributions for patients receiving pancreatic SBRT. Furthermore, the predictive accuracy of these dose models was significantly improved when dedicated models were developed for specific treatment styles. As such, despite being implemented for the same treatment site using similar dose constraints, we have shown that neural networks are capable of adapting to different treatment protocols for the

same site. This specificity is important, because one motivating force behind knowledge-based planning is to somehow quantify and characterize the essential aspects of high-quality treatment plans. Theoretically, one could obtain treatment plans from institutions boasting the best treatment outcomes, and train ANN-DMs to predict the optimal dose distributions for each new patient. In addition to promoting high-quality radiation therapy, such dose models could be used to ensure adherence to treatment protocols in large-scale clinical trials. Practically, an ANN-DM validation scheme could be implemented with relative ease, given that the model could be trained centrally by the trial coordinates and distributed to member institutions for local calculation and validation.

It should be noted that, although we have shown that ANN-DMs are capable of predicting desirable dose distributions, these models alone do not describe how to deliver these dose distributions. Yet, because these models are based off of previously delivered treatment plans, they are unlikely to propose outrageously idealistic dose distributions. ANN-DMs

FIG. 6. Voxel dose errors as a percentage of the prescribed dose are shown for Groups A and B. Top: errors from ANN-DMs are plotted alongside errors from simple linear regression models outside the PTV. Bottom: errors from ANN-DMs are plotted alongside errors from the null model inside the PTV. [Color figure can be viewed at wileyonlinelibrary.com]

trained using real plans should produce deliverable or near-deliverable dose distributions. Hence, one could envision a clinical workflow where ANN-DMs are used initially to predict a physician-approved dose distribution for each patient. With this initial head start, dosimetrists can then focus their efforts on refining each plan.

In addition, neural networks are not infallible, and should not be used to autonomously guide treatment. For example, neural networks can be susceptible to extrapolation errors whenever a new case is presented that resides outside the range of the training dataset (e.g., a PTV larger or smaller than was seen during training). Thus, as optimal plans for outlier cases accumulate, they can be used to further train the model, thereby expanding the range of values the network is proficient in. Another drawback of ANN-DM is that they do not directly calculate (and are not trained based on) DVH parameters, which are the most clinically relevant. The ANN-DMs developed here were trained according to the dose errors seen for each voxel. On the other hand, DVH errors are by their nature based off of cohorts of many voxels. However, modifications to the current approach might improve DVH

accuracy. Theoretically, one could define new input parameters that might better represent the relationship between the PTV and OARs. For example, inversely optimized arc-based treatment plans are unlikely to deliver dose through portions of the arc where the PTV and OARs are in line with one another. Ergo, as determined by the geometries of the PTV, the OARs, and the treatment machine, one could delineate regions of patient anatomy that are less likely to be irradiated. With these regions defined, a Boolean geometric parameter could be used to indicate whether or not a given voxel was in a region that was unlikely to be irradiated.

5. CONCLUSIONS

Artificial neural network dose models have been developed to reliably predict physician-approved dose distributions for pancreatic SBRT patients. The influence of different neural network features, including network complexity and input parameter choices, have been investigated and networks have been optimized accordingly. It is particularly noteworthy that significant predictive accuracy was gained by building

separate ANN-DMs for separate treatment protocols. As such, using consensus guidelines for high-quality treatment plans, neural network dose models could potentially be trained for the purposes of patient-specific plan quality validation. The inclusion of such a validation step could help promote protocol compliance, be it for a single institution, or for a large-scale clinical trial.

ACKNOWLEDGMENTS

This work was funded in part by the National Institutes of Health under award number K12CA086913, the University of Colorado Cancer Center/ACS IRG #57-001-53 from the American Cancer Society, the Boettcher Foundation, and Varian Medical Systems. These funding sources had no involvement in the study design; in the collection, analysis and interpretation of data; in the writing of the manuscript; or in the decision to submit the manuscript for publication.

^{a)} Author to whom correspondence should be addressed. Electronic mail: bernard.jones@ucdenver.edu.

REFERENCES

1. Vincent A, Herman J, Schulick R, Hruban RH, Goggins M. Pancreatic cancer. *Lancet*. 2011;378:607–620.
2. Kamisawa T, Wood LD, Itoi T, Takaori K. Pancreatic cancer. *Lancet*. 2016;388:73–85.
3. Ceha HM, van Tienhoven G, Gouma DJ, et al. Feasibility and efficacy of high dose conformal radiotherapy for patients with locally advanced pancreatic carcinoma. *Cancer*. 2000;89:2222–2229.
4. Chang DT, Schellenberg D, Shen J, et al. Stereotactic radiotherapy for unresectable adenocarcinoma of the pancreas. *Cancer*. 2009;115:665–672.
5. Koong AC, Le QT, Ho A, et al. Phase I study of stereotactic radiosurgery in patients with locally advanced pancreatic cancer. *Int J Radiat Oncol Biol Phys*. 2004;58:1017–1021.
6. Koong AC, Christofferson E, Le QT, et al. Phase II study to assess the efficacy of conventionally fractionated radiotherapy followed by a stereotactic radiosurgery boost in patients with locally advanced pancreatic cancer. *Int J Radiat Oncol Biol Phys*. 2005;63:320–323.
7. Mahadevan A, Miksad R, Goldstein M, et al. Induction gemcitabine and stereotactic body radiotherapy for locally advanced nonmetastatic pancreas cancer. *Int J Radiat Oncol Biol Phys*. 2011;81:e615–e622.
8. Chuong MD, Springett GM, Freilich JM, et al. Stereotactic body radiation therapy for locally advanced and borderline resectable pancreatic cancer is effective and well tolerated. *Int J Radiat Oncol Biol Phys*. 2013;86:516–522.
9. Timmerman R, Paulus R, Galvin J, et al. Stereotactic body radiation therapy for inoperable early stage lung cancer. *JAMA*. 2010;303:1070–1076.
10. Stinauer MA, Kavanagh BD, Scheffter TE, et al. Stereotactic body radiation therapy for melanoma and renal cell carcinoma: impact of single fraction equivalent dose on local control. *Radiat Oncol*. 2011;6:1–8.
11. Park HJ, Griffin RJ, Hui S, Levitt SH, Song CW. Radiation-induced vascular damage in tumors: implications of vascular damage in ablative hypofractionated radiotherapy (SBRT and SRS). *Radiat Res*. 2012;177:311–327.
12. Whipple AO, Parsons WB, Mullins CR. Treatment of carcinoma of the ampulla of Vater. *Ann Surg*. 1935;102:763–779.
13. Schnelldorfer T, Sarr MG. Alessandro Codivilla and the first pancreaticoduodenectomy. *Arch Surg*. 2009;144:1179–1184.
14. Are C, Dhir M, Ravipati L. History of pancreaticoduodenectomy: early misconceptions, initial milestones and the pioneers. *HPB*. 2011;13:377–384.
15. Peters JH, Carey LC. Historical review of pancreaticoduodenectomy. *Am J Surg*. 1991;161:219–225.
16. Lewis R, Drebin JA, Callery MP, et al. A contemporary analysis of survival for resected pancreatic ductal adenocarcinoma. *HPB*. 2013;15:49–60.
17. Morris PJ, Nardi GL. Pancreaticoduodenal cancer: experience from 1951 to 1960 with a look ahead and behind. *Arch Surg*. 1966;92:834–837.
18. Cameron JL, Riall TS, Coleman J, Belcher KA. One thousand consecutive pancreaticoduodenectomies. *Ann Surg*. 2006;244:10–15.
19. Birkmeyer JD, Finlayson S, Tosteson A. Effect of hospital volume on in-hospital mortality with pancreaticoduodenectomy. *Surgery*. 1999;125:250–256.
20. Abrams RA, Winter KA, Regine WF, et al. Failure to adhere to protocol specified radiation therapy guidelines was associated with decreased survival in RTOG 9704—a phase III trial of adjuvant chemotherapy and chemoradiotherapy for patients with resected adenocarcinoma of the pancreas. *Int J Radiat Oncol Biol Phys*. 2012;82:809–816.
21. Nelms BE, Robinson G, Markham J, et al. Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems. *Radiat Oncol*. 2012;2:296–305.
22. Amini A, Jones BL, Ghosh D, Scheffter TS, Goodman KA. Impact of facility volume on outcomes in patients with squamous cell carcinoma of the anal canal: analysis of the National Cancer Data Base. *Cancer*. 2017;123:228–236.
23. Shiraishi S, Moore KL. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Med Phys*. 2016;43:378–387.
24. Jabbour SK, Hashem SA, Bosch W, et al. Upper abdominal normal organ contouring guidelines and atlas: a Radiation Therapy Oncology Group consensus. *Pract Radiat Oncol*. 2014;4:82–89.