Claims:

1. A method of segmenting a volumetric image comprising a plurality of slices, the method comprising:

inputting a target slice of the volumetric image to a deep neural network (DNN) having a multi-task learning architecture, the multi-task learning architecture comprising:

a segmentation DNN that is configured to receive the target slice and output a segmentation of one or more target regions of the target slice; and

a reconstruction DNN that is configured to:

receive a plurality of adjacent slices to the target slice; and

output a reconstruction of the target slice based on the plurality of adjacent slices;

wherein the reconstruction DNN is further configured to share spatial information with the segmentation DNN, the spatial information being indicative of correlations between the adjacent slices and the target slice.

2. A method according to claim 1, wherein the reconstruction DNN comprises a convolutional feature extractor for generating first feature data from the adjacent slices, and a reconstruction downsampler for generating first reduced-dimension feature data from the first feature data at one or more scales.

3. A method according to claim 2, wherein the reconstruction DNN comprises a reconstruction upsampler for transforming the first reduced-dimension feature data to first upsampled data having the same dimensions as the first feature data.

4. A method according to claim 2 or 3, wherein the reconstruction DNN comprises one or more dimension reduction layers for applying a dimension reduction mechanism to the first feature data and/or to the first reduced-dimension feature data.

5. A method according to claim 4, wherein the dimension reduction mechanism comprises:

inputting the first feature data and/or the first reduced-dimension feature data to a 3D convolution layer;

applying an aggregation of features between adjacent slices;

applying batch normalization to the output of the 3D convolution layer; and

applying a ReLU activation function to the output of the batch normalization.

6. A method according to claim 4 or 5, wherein layers of the reconstruction downsampler are connected to layers of the reconstruction upsampler via respective ones of the dimension reduction layers by concatenation.

7. A method according to any one of claims 1 to 6, wherein the segmentation DNN comprises a convolutional feature extractor for generating second feature data from the target slice, and a segmentation downsampler for generating second reduced-dimension feature data from the second feature data at one or more scales.

8. A method according to claim 7, wherein the segmentation DNN comprises a segmentation upsampler for transforming the second reduced-dimension feature data to second upsampled data having the same dimensions as the second feature data.

9. A method according to claim 7 or 8, wherein layers of the segmentation downsampler are connected to layers of the segmentation upsampler.

10. A method according to any one of claims 7 to 9, wherein the reconstruction DNN is configured to share spatial information with the segmentation DNN by element-wise addition of output of layers of the reconstruction upsampler to output of layers of the segmentation upsampler.

11. A method according to any one of claims 1 to 10, wherein the loss function of the segmentation DNN is the 2D Intersection over Union (IoU) loss function.

12. A method according to any one of claims 1 to 11, wherein the volumetric image is a 3D medical image.

13. A method according to claim 12, wherein the 3D medical image is a 3D optical coherence tomography (OCT) image.

14. A method according to claim 13, wherein the 3D OCT image is a retinal image, and wherein the target slice corresponds to a layer of the choroid.

15. A method according to claim 14, wherein the method is repeated for a plurality of target slices, and wherein the method further comprises generating a choroidal thickness map from segmentation of the plurality of target slices.

16. A system for segmentation of a volumetric image comprising a plurality of slices, comprising:

at least one processor; and

computer-readable storage having stored thereon instructions for causing the at least one processor to carry out a method according to any one of claims 1 to 15.

17. Non-transitory computer-readable storage having instructions stored thereon for causing at least one processor to carry out a method according to any one of claims 1 to 15.