



Sequential Learning on Liver Tumor Boundary Semantics and Prognostic Biomarker Mining

Jie-Neng Chen^{1(✉)}, Ke Yan², Yu-Dong Zhang³, Youbao Tang², Xun Xu³,
Shuwen Sun³, Qiuping Liu³, Lingyun Huang⁴, Jing Xiao⁴, Alan L. Yuille¹,
Ya Zhang⁵, and Le Lu²

¹ Johns Hopkins University, Baltimore, USA

² PAII Inc, Bethesda, USA

³ The First Affiliated Hospital of Nanjing Medical University, Nanjing, China

⁴ Ping An Technology, Shenzhen, China

⁵ Shanghai Jiao Tong University, Shanghai, China

Abstract. The boundary of tumors (hepatocellular carcinoma, or HCC) contains rich semantics: capsular invasion, visibility, smoothness, folding and protuberance, etc. Capsular invasion on tumor boundary has proven to be clinically correlated with the prognostic indicator, microvascular invasion (MVI). Investigating tumor boundary semantics has tremendous clinical values. In this paper, we propose the first and novel computational framework that disentangles the task into two components: spatial vertex localization and sequential semantic classification. (1) A HCC tumor segmentor is built for tumor mask boundary extraction, followed by polar transform representing the boundary with radius and angle. Vertex generator is used to produce fixed-length boundary vertices where vertex features are sampled on the corresponding spatial locations. (2) The sampled deep vertex features with positional embedding are mapped into a sequential space and decoded by a multilayer perceptron (MLP) for semantic classification. Extensive experiments on tumor capsule semantics demonstrate the effectiveness of our framework. Mining the correlation between the boundary semantics and MVI status proves the feasibility to integrate this boundary semantics as a valid HCC prognostic biomarker.

1 Introduction

Microvascular invasion (MVI) has been clinically identified as a prognostic factor of hepatocellular carcinoma (HCC) after surgical treatment, whereas it is undetectable preoperatively on diagnostic imaging [3, 19]. Microscopic features of HCC such as tumor size, capsule and margin are hypothesized as important predictors of MVI [1, 27]. Tumor capsule, specific for hepatocarcinogenesis, was observed in 70% of progressed HCC [12]. Histologically, tumor capsule contains two layers: the inner layer is composed of tight fibrous tissue containing thin,

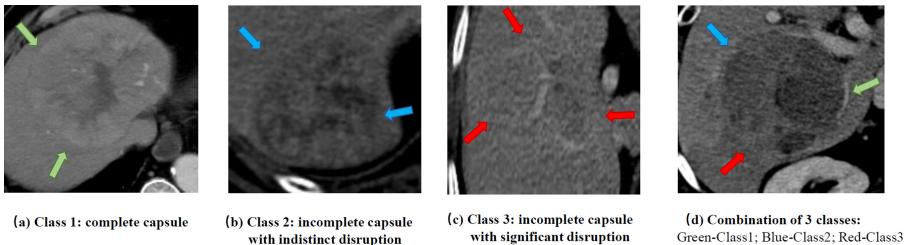


Fig. 1. Categories of tumor capsule invasion as a type of boundary semantics.

slit-like vascular channels, and the outer layer is composed of looser fibrovascular tissue [6]. Tumor capsule invasion increases the risk of vascular invasion and intrahepatic metastasis, generally indicating poor cancer patient prognosis. However, despite its strong potential, assessing the presence and integrity of radiological HCC capsule is an expert-based subjective evaluation. Inter- and intra-observer variations and lack of reproducibility are the major roadblocks limiting its wide adoption [8]. There is a critical unmet need to develop new yet effective computational methods and means to objectively and quantitatively examine the capsular invasion and focal extensional nodule to predict the precision prognosis of patients with HCC.

The capsule in radiology can be interpreted as the semantics on tumor boundary, and this problem falls into learning the tumor boundary semantics, i.e., dense classification on the boundary pixels. We propose a novel framework that disentangles the task into two pillars: the spatial localization and the sequential classification. (1) The spatial localization aims to precisely identify the vertices on tumor boundary that is extracted from predicted tumor mask. A polar coordinate transform is used to represent the boundary with radius and angle so that the boundary can be divided to N grids with equidistant angle. Then we perform an efficient vertex generator to produce the localized vertices' coordinates. (2) Our sequential learning tackles the semantic classification on the localized vertices' coordinates via sampled deep vertex features (that are concatenated with positional embedding). Finally the formed sequence features are decoded by a multilayer perceptron (MLP) for semantic classification.

To the best of our knowledge, this is the first work to solve the dense tumor boundary semantics mining problem and formulate it in a sequential learning manner. The polar coordinate transform enables us to obtain spatially uniform boundary coordinates. The vertex sequential features are sampled from multi-scale pyramid features, permitting to naturally integrate low-/mid-/high-level cues on boundary semantics. We demonstrate the effectiveness of our approach on two tumor boundary semantic datasets: capsular invasion (CAP) and focal extensional nodule (FEN). Our method improves the baseline of entangled optimization by 23.26% F1 score on CAP dataset and 10.32% F1 score on FEN dataset. Moreover, we conduct a study of prognostic biomarker mining to validate the clinical correlation between boundary semantics and MVI status.

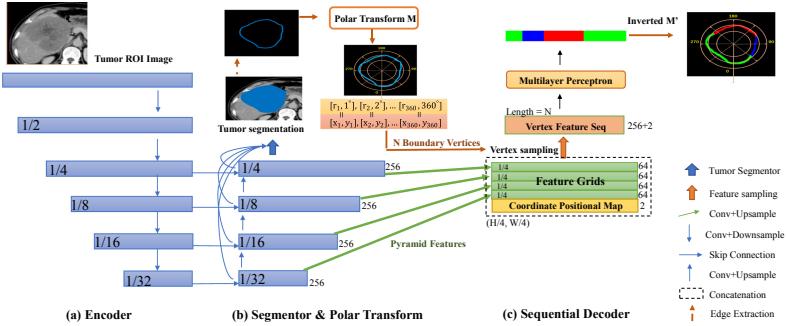


Fig. 2. Overview of our tumor boundary semantics modeling framework.

Previous Work. UPI-Net is proposed in [17] to detect semantic contour in Placental Ultrasound by binary boundary segmentation which is also studied in nature images by [2, 9, 20]. In contrast, we detect multi-classes boundary semantics independent to object semantic and/or instance category. Polar representation is used [18] to localize cell via star-convex polygons and to model foreground instances for instance segmentation [24, 25]. Sequential learning is a common task in natural language processing such as machine translation [21, 23], and various vision tasks [4, 5, 7, 16]. Moreover, to tackle with the class imbalance problem in sequential learning, Li *et al.* [13] introduced a sequential dice loss into NLP tasks as the training objective, which is adopted in our work.

2 Method

Given a tumor RoI image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$, our goal is to predict the corresponding pixel-wise label map along tumor boundary. Unlike existing approaches of directly training a segmentation network (*e.g.*, U-Net), our method converts the problem into conducting the sequential prediction on a 1D band label map $\mathbf{y} \in \mathbb{R}^{N_{angle}}$. Our overall framework is depicted in Fig. 2.

2.1 Tumor Boundary Spatial Localization

Our encoder is initialized with a ResNet-50 network [10]. ResNet features in scale $\mathbf{S} \in \{\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}\}$ will pass through the U-shape blocks, where each block consists of a $2 \times$ upsampling operator, a skip-connection, a 3×3 convolution layer, and a ReLU layer successively. Multi-scale pyramid features are generated with each scale feature $\mathbf{x} \in \mathbb{R}^{HS \times WS \times 256}$. We first build a tumor segmentation model (See Fig. 2.b) on top of the multi-scale pyramid features that upsample every scale feature into scale $1/4$ and merge them with an add operator, followed by a 3×3 convolution layer, $4 \times$ upsampling operator and a softmax activation

layer to predict the binary tumor mask $\mathbf{y}_t \in \mathbb{R}^{H \times W}$. This tumor segmentor is trained with Dice loss and cross-entropy loss. After obtaining the tumor mask, we employ the residual of the dilation and the erosion of predicted tumor mask as in [22] to generate the tumor boundary. Gaussian blurring with a 5×5 kernel is utilized to make the boundary thickness closer to the human annotation and avoid the discontinuity in the polar coordinate.

Polar coordinates is a two-dimensional coordinate system where each point on a plane is determined by a distance from a reference point and an angle from a reference direction. The reference point (analogous to the origin of a Cartesian coordinate system) is called the pole, and the ray from the pole in the reference direction is the polar axis. The distance from the pole is called the radial coordinate, and the angle is called the angular coordinate. Hence, (x_i, y_i) in Cartesian coordinate system is denoted as (r_i, θ_i) in polar coordinate system M ; the pole $(0, 0)$ in polar system is exactly the centroid of tumor mask (x_c, y_c) in Cartesian system. Cartesian position of the vertex can be recovered from the inverted transform M' , where $x_i = r_i \cos \theta_i + x_c$ and $y_i = r_i \sin \theta_i + y_c$.

We propose an efficient vertex generator to produce N boundary vertices $(x_1, y_1), \dots, (x_N, y_N)$. N rectangle grids in polar coordinate system standing for N rays with equidistant angle $\Delta\theta = \frac{360^\circ}{N}$ are generated. Specifically, grid $k \in \{1, 2, \dots, N\}$ is filled with a set of candidate vertices $(r_g, \theta_g) \in \mathbb{G}^{R \times 3}$ in polar representation, where $\theta_g \in \{(k-1)*\Delta\theta, k*\Delta\theta, (k+1)*\Delta\theta\}$, $r_g \in \{r | \theta = \theta_g\}$, and R approximates to the $\max(r | \theta = k * \Delta\theta) - \min(r | \theta = k * \Delta\theta)$. For robustness, we randomly sample a vertex point each time in N grids iteratively as a data augmentation strategy.

2.2 Deep Sequential Learning on Tumor Semantics

There are four scales of features from the feature pyramid (in Sect. 2.1) being processed to the same feature size $\frac{H}{4} \times \frac{W}{4} \times 64$ with bilinear upsampling and 3×3 convolution, and concatenated channel-wise to generate the feature $\mathbf{x}_p \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 256}$. To retain the positional information of each tumor boundary vertex and to ensure the structural relationship to be learned, we make use of the coordinate positional map [14, 15] $\mathbf{x}_{coor} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 2}$, where the channels representing the x and y Cartesian coordinates are normalized to $[-1, 1]$. Last we concatenate \mathbf{x}_p and \mathbf{x}_{coor} channel-wise to form the feature grids $\mathbf{x}_g \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 256+2}$. Given an arbitrary point (x_i, y_i) , a corresponding grid of feature $\in \mathbb{R}^{258}$ can be sampled from \mathbf{x}_g . Recall that each time the efficient vertex generator will generate a set of boundary vertices $(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$, where total N grids of feature will be sampled from \mathbf{x}_g and further sequentialized to $\mathbf{x}_{seq} \in \mathbb{R}^{N \times 258}$. A multilayer perceptron (MLP) is adopted to decode $\mathbf{x}_{seq} \in \mathbb{R}^{N \times 258}$ into 1D sequential classification prediction of $f(\mathbf{x}) \in \mathbb{R}^N$. The MLP contains two layers with a GELU non-linearity and a softmax activation. To alleviate the class imbalance where statistically the first class makes accounts for 70%, we exploit the 1D sequential dice loss [13] with cross entropy loss as the training objective for the sequential decoding. Let \mathcal{D} be the dataset and the labeled data pair

$\mathbf{x}_i, \mathbf{y}_i \in \mathcal{D}$, the sequential loss \mathcal{L}_{seq} is formulated as:

$$\mathcal{L}_{seq} = 1 - \underbrace{\sum_{\mathbf{x}_i, \mathbf{y}_i \in \mathcal{D}} 2 \frac{\sum f(\mathbf{x}_i) \mathbf{y}_i}{\sum f(\mathbf{x}_i) + \sum \mathbf{y}_i}}_{1D \text{ dice loss}} - \underbrace{\sum_{\mathbf{x}_i, \mathbf{y}_i \in \mathcal{D}} \mathbf{y}_i \cdot \log(f(\mathbf{x}_i))}_{cross \text{ entropy loss}} \quad (1)$$

For clarification, $X_p, X_{coor}, X_g, X_{seq}$ correspond to “Feature Grids”, “Coordinate Positional Map”, “the combination of Feature Grids and Coordinate Positional Map”, “Vertex Feature Sequence” in Fig. 2 respectively.

2.3 Prognostic Tumor Biomarker Mining

Microvascular invasion (MVI) score is a significant prognostic indicator of HCC in pathological imaging findings. We aim to fully exploit the correlation between MVI and our radiological imaging measurements. For each patient with multiple slices of images, we inference the data using the algorithm described above in a slice-by-slice fashion and stack the results as the prediction. The number of pixels of each class in the 3D tumor boundary is counted and divided by the total number of pixels to obtain a feature vector with three variables (e.g. [0.6, 0.1, 0.3]), where two of them are independent variables making up the patient-specific capsular biomarker. A logistic regression (LR) classifier is employed to analyze the correlation between the capsular biomarker and MVI.

3 Experiments and Discussion

Dataset Collection. A total of 364 unique patients (4049 axial slices) with pathologically confirmed liver tumor (HCC) were included in our study, which are split to 193: 58: 113 for training, validation and testing, respectively. All patients underwent standard multi-phase contrast-enhanced abdominal CT imaging within 2 weeks before surgery. *Capsular invasion (CAP)* was categorized using the following three-point scale: (1) complete capsule or invisible capsule with smooth tumor margin; (2) incomplete capsule with indistinct disruption; or (3) incomplete capsule with significant disruption. We also collect a type of tumor boundary semantics named *focal extensive nodule (FEN)* to measure the degree of protruding into the non-tumor parenchyma. FEN was assessed via a three-point criterion: (1) smooth margin without FEN; (2) slight FEN (the number of FEN is less than 3); or (3) significant FEN (the number of FEN is 3 or more). CAP and FEN classes were labeled manually using the referring three-point criteria by three board-certified radiologists on 5mm portal-venous CT images slice-by-slice. For the prognostic biomarker mining, we collected the patient-level labels of **microvascular invasion (MVI)** from their associated histopathologic examinations, where 0 stands for MVI negative; 1 for MVI positive. For inter-reader variability analysis, a subset from test-set consisting of 62 patients with 591 axial slices are labelled repeatedly by 3 radiologists Y2, Y4 and Y10, whose year-of-practice are 2, 4 and 10 respectively.

Evaluation Metrics. To evaluate the sequential classification, we perform five quantitative measures including F1 score, accuracy, AUC score, precision and recall. We adopted Dice Similarity coefficient (DSC) to evaluate the tumor segmentation accuracy that is associated with boundary extraction quality. Following [26], we perform quantitative measures including sensitivity, specificity, AUC, and Youden’s J index (J) to evaluate the logistic regression model of MVI prediction. The class weights in all LR models are adjusted to maximize Youden’s J index of (sensitivity+specificity-1).

Implementation Details. For all experiments, we apply simple data augmentations, e.g., random rotation, random resize and flipping. Models are trained with SGD optimizer with learning rate 0.01, momentum 0.9 and weight decay 1e-4. The default batch size is 64 and default training epoch is 300. The tumor segmentation branch and the sequential classification branch are optimized jointly in training. In addition, during training, we use the proposed train-time efficient vertex generator while during testing we directly generate the vertices from predicted tumor boundary.

Quantitative Comparisons. The baseline mentioned in Sect. 2 is a multi-classes segmentation network consisting of the same ResNet-50 encoder as ours and an UNet segmentation decoder. The baseline directly produces 2D semantic boundary mask. To make it comparable with the proposed method that is evaluated in sequence space, the baseline prediction mask is converted to 1D sequential band prediction by the vertex-based mask sampling, following our test-time vertex generating and sampling in Sect. 2.1. There are two branches in our framework: tumor segmentation and sequential prediction. Table 1 shows our framework can achieve 87.74% tumor DSC on the capsular invasion (CAP) dataset and 88.29% tumor DSC on the focal extensional nodule (FEN) dataset, facilitating high-quality boundary extraction and vertices generation. For sequential prediction, our framework outperforms the baseline on all measures, where the F1 score is superior by 23.25% on CAP and by 10.24% on FEN dataset. The average tumor segmentation accuracy achieves 88% DSC, indicating there is a margin of error between predicted and ground-truth tumor boundary. An upper-bound performance of our overall framework may be expected by replacing the predicted boundary with ground-truth in vertex localization. However, the quantitative results in Table 3 (*GT_Vertices*) indicate otherwise. This observation can be explained from [11] that claims the reliability of human-labelled boundary ground-truth is questionable due to ill-posed nature of boundary detection and uncertainty caused by human-annotated error. This result suggests that our framework is robust against slight shifts of boundaries due to its built-in efficient vertex generator and pyramid features. We also provide **qualitative comparison** results on the Capsular invasion prediction, as shown in Fig. 3. From there, our model is observed as the better solution to recover the boundary shape while the baseline fails sometimes (in the second row); and has stronger representation power to encode the context and distinguish the boundary semantics.

Table 1. Comparisons on CAP dataset and FEN dataset. The evaluation on sequential classification are based on F1, accuracy, AUC, precision and recall(%). We use DSC(%) to evaluate the tumor segmentation in our method.

Dataset	Model	F1	Accuracy	AUC	Precision	Recall	DSC-Tumor
CAP	Baseline	25.28	47.57	51.48	33.09	23.99	—
	Ours	48.42	64.19	64.85	54.15	48.23	87.74
FEN	Baseline	26.26	44.68	49.36	34.58	25.62	—
	Ours	36.58	52.37	59.59	42.52	37.46	88.29

Comparison with Inter-reader Variability. To quantify the inter-reader variability issue and how our approach measures against it, we compare the three radiologists' annotations (Y2, Y4, Y10) with each other and our prediction against them. The relatively poor inter-reader consistency in Table 2 show that the task is intrinsically challenging for human readers and the boundary semantics annotations lack of objectivity. Taking the ten-year practicing radiologist (Y10) as standard, our automated performance is closely equivalent to the four-year experienced radiologist (Y4), clearly proving the usability of our algorithm.

Table 2. Inter-reader variability analysis. R1 and R2 are annotation providers. Taking the ten-year practiced radiologist (Y10) as the standard, our algorithm's performance is comparable to the four-year practiced radiologist's (Y4).

R1	R2	F1	Accuracy	AUC	Precision	Recall
Y2	Y10	51.40	59.18	66.05	55.98	51.78
Y2	Y4	54.86	62.23	65.78	57.04	55.80
Y4	Y10	49.87	57.92	63.74	50.87	53.86
Average		52.04	59.78	65.19	54.63	53.81
Ours	Y2	38.78	54.16	62.52	38.00	46.28
Ours	Y4	42.49	56.67	63.15	41.70	50.20
Ours	Y10	48.60	67.65	65.24	48.72	53.26
Average		43.29	59.49	63.64	42.81	49.91

Ablation Study. We conduct the ablation study to analyze the effectiveness of different algorithm modules. The results are summarized in Table 3, where *PyraFeat* for pyramid feature, *CoordPos* for coordinate positional map, and *N_Vertices* for number of vertices/rays. (1) We investigate the impact of pyramid feature. In ablation, we only keep the $\frac{1}{4}$ scale of feature for vertex feature sampling, instead of four scales pyramid feature. From the second row of Table 3 without the pyramid feature, the model performance drops significantly on all

Table 3. Verification of the upper bound with *GT_Vertices* and the ablation study on pyramid feature, coordinate positional map, and number of vertices. All experiments are run on CAP dataset.

PyraFeat	CoordPos	N.Vertices	GT_Vertices	F1	Accuracy	AUC	Precision	Recall
✓		90		44.99	63.28	64.2	51.15	44.47
	✓	90		41.44	61.76	63.69	48.12	40.52
✓	✓	30		47.84	62.71	63.94	54.19	47.44
✓	✓	90	✓	47.5	64.23	64.79	53.63	47.23
✓	✓	90		48.42	64.19	64.85	54.15	48.23

measures. This is because the boundary semantics in our work are affected by several factors including tumor size, tumor type and the slice position, and consequently the pyramid feature in multi-scale facilitates the integrated use of low-/mid-/high level cues naturally. (2) Adding the coordinate positional map demonstrates performance robustness improvement, as the sequential feature encodes the positional dependencies. (3) Due to the vital role of boundary vertices in our framework, we analyze the impact of different numbers of vertices (i.e., the number of rays). By comparing the third row versus the fifth row in Table 3, the setting of 30 vertices yields inferior results than the default 90 vertices, implying that denser vertices are favored by our task.

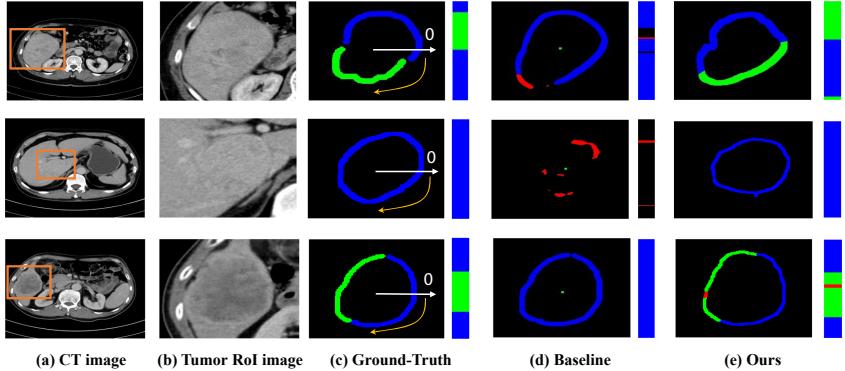


Fig. 3. Qualitative comparison. (a) Whole CT image, (b) Tumor RoI image, (c) Ground-Truth, (d) Baseline prediction, (e) Our prediction. The sequence is flatten clockwise with the starting ray of 0 angle shown in (c).

Prognostic Biomarker Mining. The inter-reader variability and the ambiguous annotation of capsular invasion in CT scans may undermine the performance since the reported results in Table 1 are far less than perfect. This motivates us to further examine the effectiveness of our system on histopathology, which are widely considered as objective observation rather than subjective interpretation.

The capsular biomarkers in the UpperBound setting uses features (pixel ratio of three classes, see Sect. 2.3) calculated from the ground-truth CAP annotations while those biomarkers in “our prediction” uses features calculated from our model’s prediction. Table 4 reports the correlation of the patient-specific capsular biomarkers and the MVI analyzed by logistic regression. It is surprising that our prediction slightly outperforms the upper bound in AUC, suggesting our predicted capsular biomarker performs comparable to radiologists’ manual annotations for the clinically important task of prognostic MVI prediction.

Table 4. The evaluated results (%) of logistic regression for predicting MVI from the capsular biomarker.

	Sensitivity	Specificity	AUC	J
UpperBound (Ground-truth)	55.88	82.85	69.36	38.73
Our prediction	67.64	72.85	70.25	40.50

4 Conclusion

Capsular invasion on tumor boundary has been clinically hypothesized of being correlated with the prognostic indicator MVI. In this paper, we present a novel quantitative computing framework on modeling the tumor boundary semantics of capsular invasion by disentangling this task to efficient spatial localization and sequential boundary semantics learning. The detected tumor boundary semantics are directly converted into a prognostic biomarker that leads to a stronger statistical correlation with MVI than the version using human annotation. For the first time we interpret the boundary semantics as an effective tumor prognostic biomarker through objective computation, and provide an alternative non-invasive way to discover the subtle sign of prognostic vascular invasion.

References

1. An, C., Kim, M.J.: Imaging features related with prognosis of hepatocellular carcinoma. *Abdominal Radiol.* **44**(2), 509–516 (2019)
2. Bertasius, G., Shi, J., Torresani, L.: High-for-low and low-for-high: efficient boundary detection from deep object features and its applications to high-level vision. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 504–512 (2015)
3. Chan, A.W., et al.: Development of pre and post-operative models to predict early recurrence of hepatocellular carcinoma after surgical resection. *J. Hepatol.* **69**(6), 1284–1293 (2018)
4. Chen, G., Chen, J., Lienen, M., Conradt, J., Röhrbein, F., Knoll, A.C.: Flgr: fixed length gists representation learning for rnn-hmm hybrid-based neuromorphic continuous gesture recognition. *Front. Neuroscience* **13**, 73 (2019)

5. Chen, J., et al.: Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint [arXiv:2102.04306](https://arxiv.org/abs/2102.04306) (2021)
6. Choi, J.Y., Lee, J.M., Sirlin, C.B.: Ct and mr imaging diagnosis and staging of hepatocellular carcinoma: part i. development, growth, and spread: key pathologic and imaging aspects. Radiology **272**(3), 635–654 (2014)
7. Dosovitskiy, A., et al.: An image is worth 16×16 words: transformers for image recognition at scale. In: ICLR (2021)
8. Ehman, E.C., et al.: Rate of observation and inter-observer agreement for li-rads major features at ct and mri in 184 pathology proven hepatocellular carcinomas. Abdominal Radiol. **41**(5), 963–969 (2016)
9. Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: 2011 International Conference on Computer Vision, pp. 991–998. IEEE (2011)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
11. Hou, X., Yuille, A., Koch, C.: Boundary detection benchmarking: beyond f-measures. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2123–2130 (2013)
12. Kojiro, M.: Histopathology of liver cancers. Best Pract. Res. Clin. Gastroenterol. **19**(1), 39–62 (2005)
13. Li, X., Sun, X., Meng, Y., Liang, J., Wu, F., Li, J.: Dice loss for data-imbalanced nlp tasks. arXiv preprint [arXiv:1911.02855](https://arxiv.org/abs/1911.02855) (2019)
14. Liang, J., Homayounfar, N., Ma, W.-C., Xiong, Y., Hu, R., Urtasun, R.: Poly-transform: deep polygon transformer for instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9131–9140 (2020)
15. Liu, R., et al.: An intriguing failing of convolutional neural networks and the coord-convolution solution. arXiv preprint [arXiv:1807.03247](https://arxiv.org/abs/1807.03247) (2018)
16. Mao, J., Xu, W., Yang, Y., Wang, J., Huang, Z., Yuille, A.: Deep captioning with multimodal recurrent neural networks (m-rnn). arXiv preprint [arXiv:1412.6632](https://arxiv.org/abs/1412.6632) (2014)
17. Qi, H., Collins, S., Alison Noble, J.: Up-net: semantic contour detection in placental ultrasound. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, p. 0 (2019)
18. Schmidt, U., Weigert, M., Broaddus, C., Myers, G.: Cell detection with star-convex polygons. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 265–273. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_30
19. Shah, S.A., et al.: Recurrence after liver resection for hepatocellular carcinoma: risk factors, treatment, and outcomes. Surgery **141**(3), 330–339 (2007)
20. Shen, W., Wang, X., Wang, Y., Bai, X., Zhang, Z.: Deepcontour: a deep convolutional feature learned by positive-sharing loss for contour detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3982–3991 (2015)
21. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. arXiv preprint [arXiv:1409.3215](https://arxiv.org/abs/1409.3215) (2014)
22. Tang, Y., Tang, Y., Zhu, Y., Xiao, J., Summers, R.M.: E2net: an edge enhanced network for accurate liver and tumor segmentation on ct scans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 512–522. Springer (2020)

23. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)
24. Xie, E., et al.: Polarmask: single shot instance segmentation with polar representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12193–12202 (2020)
25. Xu, W., Wang, H., Qi, F., Lu, C.: Explicit shape encoding for real-time instance segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 5168–5177 (2019)
26. Yao, J., Shi, Y., Lu, L., Xiao, J., Zhang, L.: DeepPrognosis: preoperative prediction of pancreatic cancer survival and surgical margin via contrast-enhanced CT imaging. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12262, pp. 272–282. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59713-9_27
27. Zhu, F., Yang, F., Li, J., Chen, W., Yang, W.: Incomplete tumor capsule on pre-operative imaging reveals microvascular invasion in hepatocellular carcinoma: a systematic review and meta-analysis. *Abdominal Radiol.* **44**(9), 3049–3057 (2019)