Flexible and Latent Structured Output Learning Application to Histology

Gustavo Carneiro¹ Tingying Peng² Christine Bayer³ Nassir Navab²

¹ ACVT, University of Adelaide, Australia **

² CAMP, Technical University of Munich, Germany

³ Department of Radiation Oncology, Technical University of Munich, Germany

Abstract. Malignant tumors that contain a high proportion of regions deprived of adequate oxygen supply (hypoxia) in areas supplied by a microvessel (i.e., a microcirculatory supply unit - MCSU) have been shown to present resistance to common cancer treatments. Given the importance of the estimation of this proportion for improving the clinical prognosis of such treatments, a manual annotation has been proposed, which uses two image modalities of the same histological specimen and produces the number and proportion of MCSUs classified as normoxia (normal oxygenation level), chronic hypoxia (limited diffusion), and acute hypoxia (transient disruptions in perfusion), but this manual annotation requires an expertise that is generally not available in clinical settings. Therefore, in this paper, we propose a new methodology that automates this annotation. The major challenge is that the training set comprises weakly labeled samples that only contains the number of MCSU types per sample, which means that we do not have the underlying structure of MCSU locations and classifications. Hence, we formulate this problem as a latent structured output learning that minimizes a high order loss function based on the number of MCSU types, where the underlying MCSU structure is flexible in terms of number of nodes and connections. Using a database of 89 pairs of weakly annotated images (from eight tumors), we show that our methodology produces highly correlated number and proportion of MCSU types compared to the manual annotations.

Keywords: Weakly Supervised Training, Latent Structured Output Learning, High order loss function

1 Introduction

The majority of human tumours contain chronic (limitations in oxygen diffusion) and acute (local disturbances in perfusion) hypoxic regions, which lead to poor clinical prognosis in treatments based on radiotherapy and chemotherapy [1]. While chronic hypoxia (CH) promotes the death of normal and tumor cells [2], acute hypoxia (AH) leads to tumor aggressiveness, so it is important to estimate the number and proportion of hypoxic regions in tumors to improve the clinical prognosis of such treatments [1]. Matei et al. [2] proposed a manual annotation

^{**} Gustavo Carneiro thanks the Alexander von Humboldt Foundation (Fellowship for Experienced Researchers). This work was partially supported by the Australian Research Council Projects funding scheme (project DP140102794).



Fig. 1: Manual annotation using the HE and IF images as inputs and producing a count and proportion of the MCSU types present in the histological specimen.

of the number and proportion of hypoxic regions using (immuno-)fluorescence (IF) and hematoxylin and eosin (HE) stained images of a histological specimen, involving the following steps (Fig. 1): 1) registration of the IF and HE images; 2) delineation of the vital tumor region; 3) detection of microcirculatory supply units (MCSU), which are areas supplied by microvessels; 4) classification of MCSUs into normoxia (N - normal oxygenation supply), CH or AH; and 5) computation of the number and proportion of MCSU types. This annotation requires expertise that is generally not available in clinical settings, which makes it a good candidate for automation. A major hurdle is that this annotation [2] contains only the final number and proportion of MCSU types, without indication of MCSU locations, sizes and labels (an MCSU has size of around $200 \mu m$ and class appearance as defined in Fig. 2). Therefore, this is a weakly supervised and multi-class structured learning problem that is formulated in this paper as a latent structured output problem [3] that minimizes a high order loss function [6] based on the mismatch between the manual and automated estimation of the number of MCSU types, where this latent structure is flexible in terms of the connections and number of MCSUs.

Literature Review: Although new, our problem is similar the segmentation of brain structures [7–9], involving a detection of sparse structures and multiclass classification. However, different from our problem, the segmentation of brain structures is formulated as a strongly supervised problem, where it is possible to use the position and shape priors. The detection of lymph nodes [10] also deals with the identification of sparse structures without priors. In contrast to our problem, lymph node detection is strongly supervised and concerns a binary classification problem. The automated detection and localization of multiple organs [11, 12] also deals with sparse detection and multi-class classification, but it is strongly supervised and one can use position and shape priors. There are a few problems formulated as weakly supervised latent structured output learning [13–15], but they present



Fig. 2: MCSU classes appearance [2].

some differences compared to our problem, as detailed below. The tracking of

indistinguishable translucent objects [13] uses a stronger and lower level annotation, consisting of the identification of the objects before and after occlusion. In the semantic segmentation [14, 15], images are annotated with a set of classes, where the pixel-level annotation is not available, but these methodologies use lower level loss functions and deal with non-sparse segmentation problems.

Contributions: Our contribution is a new weakly supervised latent structured output learning methodology for the detection and multi-class classification of sparse structures in multimodal cytological images that is trained with the minimization of a high order loss function, where the main novelty is the flexible structure of the latent MCSU structure in terms of number of nodes and connections. In addition, this is the first methodology for the automated classification of oxygenation levels in multimodal cytological images. We analised a database of 89 pairs of IF and HE images (from eight tumors), where 16 pairs of images from two tumors were for training local MCSU multi-class classifiers, and 73 pairs of images from six tumors were for training and testing the latent structured output learning methodology. Using a leave-one-tumor-out cross validation experiment, we obtain a high correlation between the manual and automated annotations in terms of the number and proportion of MCSU types.

2 Methodology

Our methodology depends on a dataset $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{v}_n, \mathbf{y}_n)\}_{n=1}^N$, where $\mathbf{x} = \{\mathbf{x}^{(\mathrm{IF})}, \mathbf{x}^{(\mathrm{HE})}\}$ is the input IF and HE images, with $\mathbf{x}^{(\mathrm{IF})}, \mathbf{x}^{(\mathrm{HE})} : \Omega \to \mathbb{R} \ (\Omega \in \mathbb{R}^2$ denotes the image lattice), $\mathbf{v} : \Omega \to \{0, 1\}$ is the vital tumor mask (Fig. 1), and $\mathbf{y} \in \mathcal{Y} \subseteq \mathbb{N}^3$ denotes the annotation of the number of normotic (N), chronic hypoxic (CH) and acute hypoxic (AH) MCSUs. The hidden structure is represented by the graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where the nodes in \mathcal{V} denote the MCSUs and edges in \mathcal{E} represent their connections, and each node is associated with a label $c \in \{1, 2, 3, 4\}$ (1 stands for N, 2 for CH, 3 for AH and 4 for Necrosis). We include the class Necrosis (Ne) because the vital tumor mask \mathbf{v} often includes necrotic regions that must be processed during learning and inference. The structure and labeling of this graph are formed by an algorithm parametrized by the latent variable $h \in \mathcal{H}$ and the output variable \mathbf{y} , as described below (see Fig. 3).

2.1 Inference and Learning

We formulate our problem as a latent structured support vector machine parameterized by \mathbf{w} , where the inference optimizes the following objective function:

$$(\mathbf{y}^*, h^*) = \arg \max_{\mathbf{y} \in \mathcal{Y}, h \in \mathcal{H}} \mathbf{w}^\top \Psi(\mathbf{x}, \mathbf{y}, h).$$
(1)

In (1), we have $\Psi(\mathbf{x}, \mathbf{y}, h) = [f_1^{(1,1)}, ..., f_4^{(1,1)}, ..., f_1^{(1,K)}, ..., f_4^{(1,K)}, f^{(2,1)}, ..., f^{(2,L)}]$, where $f_c^{(1,k)} = \sum_{v \in \mathcal{V}} \delta(m_v(\mathbf{y}) - c)\phi^{(1,k)}(c, \mathbf{x}; \theta^{(1,k)})$ with $m_v(\mathbf{y}) \in \{1, 2, 3, 4\}$ denoting the label of node $v \in \mathcal{V}$ estimated with \mathbf{y} as described below in (3), $\delta(.)$ is the Dirac delta function and $k \in \{1, ..., K\}$ with $\phi^{(1,k)}(c, \mathbf{x}; \theta^{(1,k)}) = -\log P^{(k)}(c|\mathbf{x}_v, \theta^{(1,k)})$ representing the k^{th} unary potential function defined below in (3) and representing the negative log probability of assigning class c to node v, and $f^{(2,l)} = \sum_{(v,t)\in\mathcal{E}} \phi^{(2,l)}(c_v, c_t, \mathbf{x}; \theta^{(2,l)})$ for $l \in \{1, ..., L\}$ with
$$\begin{split} \phi^{(2,1)}(c_v,c_t,\mathbf{x};\theta^{(2,l)}) &= (1-\delta(c_v-c_t))g(c_v,c_t,\mathbf{x};\theta^{(2,l)}) \text{ representing the binary} \\ \text{potential function that measures the compatibility (indicated by <math>g(.)$$
) between nodes v and t when their labels are different (indicated by the Dirac $\delta(.)$). We consider the following binary potentials: 1) $g(c_v,c_t,\mathbf{x};\theta^{(2,1)}) = 1/||\mathbf{p}_v - \mathbf{p}_t||$ (where $\mathbf{p}_v \in \Omega$ denotes the position of node v in the image), 2) $g(c_v,c_t,\mathbf{x};\theta^{(2,2)}) = 1/||\mathbf{r}_v - \mathbf{r}_t||$ (where $\mathbf{r}_v = [P^{(k)}(c_v|\mathbf{x},\theta^{(1,k)})]_{c_v \in \{1,\ldots,4\},k \in \{1,\ldots,K\}} \in \mathbb{R}^{4 \times K}$ is a vector of the classifier responses for each class in node v); and 3) $g(c_v,c_t,\mathbf{x};\theta^{(2,3)}) = 1/(||\mathbf{p}_v - \mathbf{p}_t|| \times ||\mathbf{r}_v - \mathbf{r}_t||)$.

The learning procedure is formulated as [4]:

ŝ

$$\begin{array}{ll} \underset{\mathbf{v},\{\xi_n\}_{n=1}^N}{\text{minimize}} & \frac{1}{2} \|\mathbf{w}\|^2 + \frac{C}{N} \sum_{n=1}^N \xi_n \\ \text{subject to} & \underset{h_n \in \mathcal{H}}{\max} \mathbf{w}^\top \Psi(\mathbf{x}_n, \mathbf{y}_n, h_n) - \mathbf{w}^\top \Psi(\mathbf{x}_n, \widehat{\mathbf{y}}_n, \widehat{h}_n) \ge \Delta(\mathbf{y}_n, \widehat{\mathbf{y}}_n) - \xi_n \\ & \xi_n \ge 0, \forall \widehat{\mathbf{y}}_n \in \mathcal{Y}, \forall \widehat{h}_n \in \mathcal{H}, n = 1, ..., N, \end{array} \tag{2}$$

where ξ_n are the slack variables and $\Delta(\mathbf{y}_n, \hat{\mathbf{y}}_n) = \sum_{c=1}^3 |\mathbf{y}_n(c) - \hat{\mathbf{y}}_n(c)|$ measures the loss between annotations \mathbf{y}_n and $\hat{\mathbf{y}}_n$. The problem in (2) is solved by the following concave-convex procedure [16]: 1) estimation of the latent variable value consistent with annotations and current estimate for \mathbf{w} , as in $\max_{h_n \in \mathcal{H}} \mathbf{w}^\top \Psi(\mathbf{x}_n, \mathbf{y}_n, h_n)$; and 2) new estimation of \mathbf{w} using (2) given $\{h_n\}_{n=1}^N$ from step 1. Note that the estimation of \mathbf{w} is based on the cutting plane algorithm [17] that iteratively solves a loss augmented inference problem that inserts a new constraint in the set of most violated constraints with $(\hat{\mathbf{y}}_n, \hat{h}_n) =$ arg $\max_{\mathbf{y} \in \mathcal{Y}, h \in \mathcal{H}} \Delta(\mathbf{y}_n, \mathbf{y}) + \mathbf{w}^\top \Psi(\mathbf{x}, \mathbf{y}, h)$. The inference used for this loss augmented problem and (1) is based on graph cuts (GC) [18], where the high order loss function $\Delta(\mathbf{y}_n, \hat{\mathbf{y}}_n)$ is integrated into GC based on the decomposition of [6].

2.2 The Flexible and Latent Structure $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

The flexible and latent structure represented by the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is needed to built $\Psi(.)$ in (1) and (2). The estimation of \mathcal{G} starts with the detection and classification of microvessel pixels (leftmost image in Fig. 3-(b)) using the IF and HE images (Fig. 3-(a)). We define a variable $\mathbf{t}: \Omega \to \{0, 1\}$, where $\mathbf{t}_i = 1$ if the red channel of the IF image at $i \in \Omega$ is larger than $\tau = 0.1$ (from the range [0,1], otherwise $\mathbf{t}_i = 0$ (see yellow dots in the first image of Fig. 3-(b)). Using the sketches in Fig. 2, we annotate image samples for training the following multi-class classifiers: 1) Adaboost [19], 2) linear SVM [20], 3) random forest [21], and 4) deep convolutional neural networks [22]. The features used by these classifiers are composed of the pixel values extracted from a patch \mathbf{x}_i of size $200\mu m$, where $\mathbf{t}_i = 1$. Note that, as explained above, the class Ne must be added, where in general, a necrotic patch comprises a red center with black pixels around it in the IF image and a dark purple color in the HE image. This process results in four classifiers: $\{P^{(k)}(c|\mathbf{x}_i, \theta^{(1,k)})\}_{k=1}^K$ (with K = 4). We show the results from a majority voting process of the four classifiers in the middle image of Fig. 3-(b). We can then form an initial graph using the microvessel pixels, with $\mathcal{G}^{ini} = (\mathcal{V}^{ini}, \mathcal{E}^{ini})$, with $\mathcal{V}^{ini} = \{i | \mathbf{t}_i = 1\}$, and the edges \mathcal{E}^{ini} defined by Delaunay triangulation (rightmost image in Fig. 3-(b)).



Fig. 3: Building and labeling \mathcal{G} . From the IF/HE images (a), microvessel pixels are detected, classified and structured in an initial graph (b), which is modified to represent the MCSU structure (c) that are then used to form $\Psi(.)$ for (1) and (2).

The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is built and labeled using h and \mathbf{y} , as follows: 1) estimate the graph structure by running a minimum spanning tree clustering algorithm [23] over the graph \mathcal{G}^{ini} , where the edge weight between nodes i and j is defined as $\|\mathbf{p}_i - \mathbf{p}_j\| \times \|\mathbf{r}_i - \mathbf{r}_j\|$ (\mathbf{p}_i is the 2-D position of i and \mathbf{r}_i is the response from the classifiers - this emphasizes that nearby microvessel pixels that have similar classifier responses must be in the same MCSU), and each cluster \mathcal{C} must have a diameter smaller than $h \times 200 \mu m$, with $h \in [0.5, 2]$ (this diameter is measured by $\max_{i,j\in\mathcal{C}} \|\mathbf{p}_i - \mathbf{p}_j\|$ with $\mathcal{C} = \{i|i \in \mathcal{V}^{ini}\}$ denoting the set of \mathcal{G}^{ini} nodes belonging to the same cluster); and 2) assuming $\{\mathcal{C}_v\}_{v=1}^{|\mathcal{V}|}$ are the clusters for the nodes $v \in \mathcal{V}$ from step (1), the graph labeling (to be used by $\Psi(.)$ in (1) and (2)) uses the annotation \mathbf{y} as follows:

$$\begin{array}{ll} \underset{\mathbf{M}}{\operatorname{minimize}} & -\|\mathbf{M} \odot \mathbf{P}\|_{F}^{2} + \sum_{c=1}^{3} \left(\mathbf{y}(c) - \|\mathbf{M} \odot \mathbf{E}_{c}\|_{F}^{2} \right)^{2} \\ \text{subject to} & \mathbf{1}_{4}^{\top} \mathbf{M} = \mathbf{1}_{|\mathcal{V}|}^{\top}, \ \mathbf{M} \in \{0, 1\}^{4 \times |\mathcal{V}|}, \end{array}$$

$$(3)$$

where $\mathbf{P} \in \mathbb{R}^{4 \times |\mathcal{V}|}$, with $\mathbf{P}(c, v) = \prod_{k=1}^{K} P^{(k)}(c|\mathbf{x}_{v}, \theta^{(1,k)})$ for $c \in \{1, 2, 3, 4\}$ and $v \in \mathcal{V}$ (note that $P^{(k)}(c|\mathbf{x}_{v}, \theta^{(1,k)}) = \prod_{i \in \mathcal{C}_{v}} P^{(k)}(c|\mathbf{x}_{i}, \theta^{(1,k)})$ in (1)), $\mathbf{E}_{1} = [\mathbf{1}_{|\mathcal{V}|}, \mathbf{0}_{|\mathcal{V}|}, \mathbf{0}_{|\mathcal{V}|}]^{\top} \in \{0, 1\}^{4 \times |\mathcal{V}|}$ denotes a matrix with ones in first row and zeros elsewhere (similarly for c = 2, 3 with ones in rows 2 and 3), $\mathbf{1}_{N}$ and $\mathbf{0}_{N}$ represent a size N column vector of ones or zeros, $\|.\|_{F}$ denotes the Frobenius norm, \odot represents the Hadamard product, and the summation varies from 1 to 3 because \mathbf{y} has the annotation for three classes only. The optimization in (3) minimizes the objective function by maximizing the label assignment probability and minimizing the difference between the number of MCSU classes in \mathbf{M} and in the variable \mathbf{y} . We relax the second constraint to $\mathbf{M} \in [0, 1]$ to make the original integer programming problem feasible. The edge set \mathcal{E} is obtained with Delaunay triangulation (left in Fig. 3-(c)). Note that \mathbf{M} in (3) contains the label of each node $v \in \mathcal{V}$ needed in (1), with $m_v(\mathbf{y}) = \arg \max_{c \in \{1,...,4\}} \mathbf{M}(c, v)$ (right in Fig. 3-(c)). The discrepancy in the number of microvessel pixels and MCSUs shown in Fig. 3(b)-(c) is due to the fact that MCSUs are formed by a set of microvessel pixels, where MCSUs could have been cut in different directions (parallel, oblique, or transversal) during material preparation, and also because MCSUs can vary in size.

3 Materials and Methods

We use the material from [2], consisting of five xenografted human squamous cell carcinoma lines of the head and neck (FaDu), which were transplanted subcutaneously into the right hind leg of nude mice. Each whole tumor cryosection was scanned and photographed using AxioVision 4.7 and the multidimensional and mosaix modules. The IF images of tumor cryosections were prepared with three separate stainings: Pimonidazole was used for visualizing hypoxia (green regions), CD31 for microvessels (red regions), and Hoechst 33342 for perfusion stain (blue). Next, the cover slip was removed to stain the same slice with HE. In total, there are 89 pairs of IF and HE images from eight tumors, where 16 pairs from two tumors are used for training the classifiers $\{P^{(k)}(c|\mathbf{x}_i, \theta^{(1,k)})\}_{k=1}^4$, and the remaining 73 pairs from six tumors for training and testing our proposed weakly supervised latent structured output learning methodology. We run a sixfold cross validation experiment, leaving one tumor out in each run. These pairs of IF and HE images are registered [24] and downsampled such that the resolution is approximately $10\mu m$ per pixel, and a manually delineated mask is used to mark the vital tumor tissue (Fig. 1).

The estimation of \mathbf{y}^* and \hat{h}^* in (1) and the loss augmented inference in (2) to estimate $\hat{\mathbf{y}}_n$ and \hat{h}_n are based on graph cut (alpha-expansion) [18] with $\mathcal{H} = \{0.5, 1, 1.5, 2\}$. We compare our results with an ideal method based on an observed structured SVM, where h is set with a value that best approximates the manually annotated number of MCSUs (that is, h is treated as observed in this ideal method), as in $h^* = \arg \min_h ||\mathbf{1}^\top \mathbf{y} - |\mathcal{V}|||$, where $|\mathcal{V}|$ is the number of MCSUs in \mathcal{G} (Sec. 2.2). This result produced by this ideal method can be seen as the best case scenario for our latent structured output learning problem. Finally, the quantitative experiment measures the correlation of the percentage and counting of MCSU types between manual and automated annotations in the test sets, using Bland Altman plots [25], which shows the number of samples, sum of squared error (SSE), Pearson r-value squared, linear regression, and p-value.

4 Results

Figure 4 (a-b) shows the Bland Altman plots for the proposed methodology for the percentage and counting of MCSU classes, which can be compared to the ideal method in Fig. 4(c-d). In all cases in Fig. 4, the correlation coefficient (r^2) is around 0.8 and above with p-values significantly smaller than 0.01, showing strong correlation results. Finally, Fig. 5 shows examples (using different tumors from the test set) of the annotations produced by our proposed methodology.

5 Discussion and Conclusion

Fig. 4 shows that the proposed methodology produces a result that is comparable to the ideal method (with a "known" h) in the case of percentage of MCSU



Fig. 4: Bland Altman graphs of the percentage (left) and counting (right) results for the proposed methodology (top) and the ideal method with observed h (bottom).

classes. In the case of counting the MCSU types, our methodology presents a larger variance, but a similar bias. Nevertheless, the correlation coefficient obtained for both cases (MCSU percentage and counting) is large, with values around 0.8 and above and p-values << 0.01, indicating strong correlation results. The qualitative results in Fig. 5 show that our final classification results are similar to the ground truth annotation. These results provide evidence that our approach is potentially useful in a clinical setting for the automated annotation of oxygenation levels in multimodal cytological images.

References

- 1. Bayer, C., Vaupel, P.: Acute versus chronic hypoxia in tumors. Strahlentherapie und Onkologie **188**(7) (2012) 616–627
- Maftei, C.A., et al.: Changes in the fraction of total hypoxia and hypoxia subtypes in human squamous cell carcinomas upon fractionated irradiation: evaluation using pattern recognition in microcirculatory supply units. Radiotherapy and Oncology 101(1) (2011) 209–216
- 3. Yu, C.N.J., Joachims, T.: Learning structural svms with latent variables. In: ICML (2009) 1169–1176
- 4. Kumar, M.P.: Weakly Supervised Learning for Structured Output Prediction. PhD thesis, Ecole Normale Supérieure de Cachan (2014)
- Lou, X., Hamprecht, F.: Structured learning from partial annotations. arXiv preprint arXiv:1206.6421 (2012)
- Pletscher, P., Kohli, P.: Learning low-order models for enforcing high-order statistics. In: AISTATS (2012) 886–894
- 7. Carneiro, G. et al.: Semantic-based indexing of fetal anatomies from 3-d ultrasound data using global/semi-local context and sequential sampling. In: CVPR. (2008)



Fig. 5: Automated annotations produced by the proposed methodology. Please see additional results on the supplementary material.

- 8. Patenaude, B. et al.: A bayesian model of shape and appearance for subcortical brain segmentation. Neuroimage 56(3) (2011) 907–922
- 9. Tu, Z. et al.: Brain an atomical structure segmentation by hybrid discriminative/generative models. TMI ${\bf 27}(4)~(2008)~495{-}508$
- 10. Barbu, A. et al.: Automatic detection and segmentation of lymph nodes from ct data. TMI ${\bf 31}(2)~(2012)~240{-}250$
- 11. Pauly, O. et al.: Fast multiple organ detection and localization in whole-body mr dixon sequences. In: MICCAI (2011).
- Zhou, S.K.: Discriminative anatomy detection: Classification vs regression. Pattern Recognition Letters 43 (2014) 25–38
- 13. Fiaschi, L. et al.: Tracking indistinguishable translucent objects over time using weakly supervised structured learning. In: CVPR (2014).
- 14. Mahapatra, D. et al.: Weakly supervised semantic segmentation of crohn's disease tissues from abdominal mri. In: ISBI (2013).
- Quellec, G. et al.: Weakly supervised classification of medical images. In: ISBI (2012)
- Yuille, A.L., Rangarajan, A.: The concave-convex procedure. Neural computation 15(4) (2003) 915–936
- Joachims, T. et al.: Cutting-plane training of structural syms. Machine Learning 77(1) (2009) 27–59
- Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. TPAMI 23(11) (2001) 1222–1239
- 19. Zhu, J. et al.: Multi-class adaboost. Statistics and Its (2009)
- Tsochantaridis, I. et al.: Support vector machine learning for interdependent and structured output spaces. In: ICML (2004).
- 21. Breiman, L.: Random forests. Machine learning 45(1) (2001) 5–32
- 22. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In:NIPS (2012)
- 23. Grygorash, O., Zhou, Y., Jorgensen, Z.: Minimum spanning tree based clustering algorithms. In: ICTAI (2006).
- Peng, T. et al.: Deformable registration of multi-modal microscopic images using a pyramidal interactive registration-learning methodology. In: Biomedical Image Registration. Springer (2014) 144–153
- 25. Altman, D.G., Bland, J.M.: Measurement in medicine: the analysis of method comparison studies. The statistician (1983) 307–317