Linear and Non-Linear Dimensionality Reduction for Hyperspectral Overlapping Chromosome Segmentation

Najoua Zeffate[1⋆], Mohammed El Amine Bechar[1], Jean-Marie Guyader[1], Nathalie Douet-Guilbert[2,3], Marwa El Bouz[1⋆⋆], and Marie-Bérengère Troadec[2,3⋆⋆]

[1] LabISEN, ISEN Ouest, Brest, France
[2] Univ Brest, Inserm, EFS, UMR 1078, GGB, Brest, France
[3] CHRU Brest, Service de génétique, Laboratoire de génétique chromosomique, Brest, France
najoua.zeffate@isen-ouest.yncrea.fr

Abstract. Hyperspectral imaging (HSI) provides rich spectral information for biomedical image analysis but exhibits high dimensionality, which increases computational complexity and introduces redundant information. This paper investigates the impact of dimensionality reduction on hyperspectral overlapping chromosome segmentation . Three linear methods, namely Principal Component Analysis (PCA), Non-negative Matrix Factorization (NMF), and Factor Analysis (FA), and three nonlinear methods, namely Kernel Principal Component Analysis (Kernel PCA), t-distributed Stochastic Neighbor Embedding (t-SNE), and Isomap, are evaluated. Each method projects the original hyperspectral cube into a three-component representation used as input to 2D U-Net architecture. The resulting models are compared with a 2D U-Net using RGB images, as well as 2D and 3D U-Nets operating on the complete hyperspectral cube. Experiments are conducted using five-fold cross-validation and evaluated using the Dice coefficient and Intersection over Union (IoU). Paired Wilcoxon signed-rank tests are performed to assess statistical significance. The results show that dimensionality reduction significantly improves segmentation performance compared with both RGB images and the complete HSI cube. Although Isomap achieved the highest average Dice and IoU scores, pairwise statistical analyses revealed no significant differences between most dimensionality reduction methods. These findings demonstrate that dimensionality reduction is an effective preprocessing strategy for hyperspectral overlapping chromosome segmentation.

Keywords: Hyperspectral imaging · Overlapping chromosome · Segmentation · Dimensionality reduction

⋆ Corresponding author
⋆⋆ Co-last authors

1 Introduction

Chromosome analysis is an important task in cytogenetics, where accurate chromosome segmentation is a key step for downstream analysis such as karyotyping and abnormality detection [16]. In particular, overlapping chromosomes represent a challenging segmentation problem, as they often exhibit ambiguous boundaries and similar visual appearances. These difficulties can lead to inaccurate segmentation masks, especially when relying only on conventional intensity or RGB information [21]. Hyperspectral imaging (HSI) has emerged as a promising imaging modality for biomedical image analysis, as it captures both spatial and spectral information over a large number of narrow wavelength bands [7]. Compared with conventional RGB imaging, HSI provides richer spectral signatures that can help distinguish materials or biological structures with similar visual appearance but different spectral responses [7]. However, the high dimensionality of hyperspectral data also introduces several challenges, including increased computational cost, spectral redundancy, and a higher risk of overfitting when training deep learning models on limited annotated datasets [22]. Dimensionality reduction methods have been widely used to extract compact representations from hyperspectral data [9]. Linear methods, such as Principal Component Analysis (PCA) [4], Non-negative Matrix Factorization (NMF)[6], and Factor Analysis (FA)[13], aim to project the data into a lower-dimensional subspace while preserving dominant spectral variability or latent factors [2]. Nonlinear methods, such as Kernel PCA [11], t-distributed Stochastic Neighbor Embedding (t-SNE)[8], and Isomap[14], can capture more complex structures in the data by considering nonlinear relationships or manifold geometry [15, 5]. However, it remains unclear which type of dimensionality reduction is most suitable for hyperspectral overlapping chromosome segmentation . For hyperspectral data, three main strategies can be considered. The first consists of directly processing the full hyperspectral cube with a 2D convolutional network by treating the spectral bands as input channels[12]. The second relies on a 3D convolutional network to jointly exploit spatial and spectral information [1, 18]. However, directly processing the full spectral information may increase computational complexity and include redundant information. The third strategy consists of reducing the spectral dimensionality of the hyperspectral cube before segmentation, allowing the use of standard 2D convolutional networks while retaining informative spectral content [19, 3]. In this work, we investigate the impact of linear and nonlinear dimensionality reduction techniques on the segmentation of chromosome crossings from hyperspectral images. Each method reduces the original hyperspectral cube into a three-component representation, which is then used as input to 2D U-Net architecture [10]. The resulting segmentation performance is compared with two baseline approaches: a 2D U-Net trained separately on RGB images and on the full hyperspectral cube, and a 3D U-Net trained directly on the full hyperspectral cube. The contributions of this study are threefold: (i) we provide a systematic comparison of representative linear and nonlinear dimensionality reduction methods for hyperspectral overlapping chromosome segmentation ; (ii) we evaluate whether compact three-

component hyperspectral representations improve segmentation compared with RGB and full HSI inputs; and (iii) we analyze the statistical significance of the observed performance differences using paired non-parametric tests.

2 Method

2.1 Overview

The proposed framework aims to evaluate the effect of dimensionality reduction on hyperspectral overlapping chromosome segmentation . Given a hyperspectral image cube, the spectral dimension is first reduced to three components using a selected dimensionality reduction method. The resulting three-channel representation is then used as input to a 2D U-Net segmentation network. This setting a fair comparison between different reduced representations, since the same segmentation architecture and training protocol are used for all dimensionality reduction methods.

Let $\mathbf{X} \in \mathbb{R}^{H\times W\times B}$ denote a hyperspectral image, where H and W are the spatial dimensions and B is the number of spectral bands. In our experiments, each cube has a spatial size of 64×64 pixels and contains 104 spectral bands. For dimensionality reduction, each cube is reshaped into a two-dimensional matrix $\mathbf{X}' \in \mathbb{R}^{N\times B}$, where $N = H \times W$ is the number of pixels. Each pixel is therefore represented by a spectral vector of length B. A dimensionality reduction function $f(\cdot)$ is then applied to obtain a compact representation: $\mathbf{Z} = f(\mathbf{X}'), \quad \mathbf{Z} \in \mathbb{R}^{N\times 3}$. The reduced matrix $\mathbf{Z}$ is finally reshaped into a three-channel image $\mathbf{Z}' \in \mathbb{R}^{H\times W\times 3}$, which is used as input to the 2D U-Net.

2.2 Dimensionality Reduction

Six dimensionality reduction methods are evaluated in this study. They are divided into two groups: linear and nonlinear approaches.

The linear methods include PCA, NMF, and FA. PCA projects the data onto orthogonal directions that maximize the variance of the spectral data [4]. NMF decomposes the hyperspectral data into non-negative components, which can provide an additive parts-based representation [6]. FA models the observed spectral bands using a smaller number of latent factors and a noise term [13].

The nonlinear methods include Kernel PCA with an RBF kernel (KPCA-RBF), t-SNE, and Isomap. Kernel PCA extends PCA by implicitly mapping the input data into a higher-dimensional feature space using a kernel function [11]. t-SNE aims to preserve local neighborhood relationships by matching probability distributions between the high-dimensional and low-dimensional spaces [8]. Isomap constructs a neighborhood graph and seeks to preserve geodesic distances along the underlying data manifold [14].

All dimensionality reduction methods are constrained to produce three components. This choice is motivated by two considerations: First, it enables direct compatibility with standard 2D U-Net architectures designed for three-channel

inputs. Second, this choice is supported by previous work showing that most of the spectral variance of hyperspectral chromosome images can be represented by the first three principal components [20].

2.3 Segmentation Networks

For all reduced hyperspectral representations, segmentation is performed using a 2D U-Net architecture [10]. The network takes a three-channel image as input and outputs a binary segmentation mask. Using the same segmentation architecture across all dimensionality reduction methods ensures that performance differences are mainly attributable to the input representation rather than to changes in the segmentation model.

Two baseline approaches are also considered. The first uses a 2D U-Net architecture trained separately on RGB images and on the full 104-band hyperspectral cube. This allows the effect of dimensionality reduction to be evaluated while keeping the segmentation architecture unchanged. The second baseline is a 3D U-Net trained directly on the full hyperspectral cube to jointly exploit its spatial and spectral dimensions.

2.4 Training and Evaluation Protocol

The experiments were conducted on a private dataset containing 266 hyperspectral cubes of overlapping chromosomes, each with a spatial resolution of 64×64 pixels and 104 spectral bands. A fixed subset of 60 cubes was reserved for testing. The remaining 206 samples were divided according to a five-fold cross-validation protocol, with 165 training and 41 validation samples per fold, except for the first fold, which contained 164 and 42 samples, respectively. The same data partitions and independent test set were used across all experiments. All segmentation models were trained for 500 epochs with a batch size of 8. The 2D U-Net was optimized using Adam with an initial learning rate of 2×10^{-4} and binary cross-entropy loss. Its learning rate was linearly decreased during the final training epochs. The 3D U-Net was trained using Adam with the same initial learning rate, a weight decay of 10^{-5}, and cross-entropy loss. All experiments were conducted under Linux on an NVIDIA TITAN V GPU equipped with 12 GB of HBM2 memory. Segmentation performance is assessed using the Dice coefficient and Intersection over Union (IoU). To assess whether the observed differences are statistically significant, paired Wilcoxon signed-rank tests [17] are performed on the IoU scores of the same 60 independent test samples ($n = 60$) . First, each dimensionality reduction method is compared with the RGB and full HSI baselines. Then, pairwise comparisons are performed between the dimensionality reduction methods to evaluate whether one reduced representation provides a statistically significant advantage over the others. For the main experiments, all dimensionality-reduction methods were configured to produce three components. The remaining hyperparameters were kept at their default values, with an RBF kernel used for Kernel PCA. Specifically, t-SNE used perplexity = 30, Isomap used $n_{\text{neighbors}} = 5$, and Kernel PCA used $\gamma = 1/104$. A sensitivity analysis was

conducted for the main hyperparameters of the nonlinear methods. The t-SNE perplexity and Isomap neighborhood size were varied over $\{5, 10, 15, 30, 50\}$. For Kernel PCA, γ was evaluated over $\{\gamma_0/4, \gamma_0/2, \gamma_0, 2\gamma_0, 4\gamma_0\}$, where $\gamma_0 = 1/104$. All other parameters were kept fixed.

3 Results and Discussion

3.1 Quantitative and qualitative Results

Table 1: Quantitative comparison of linear and nonlinear dimensionality reduction methods against RGB and HSI baselines for overlapping chromosome segmentation.

	Input	**Model (U-Net)**	**IoU (%) ↑**	**Dice (%) ↑**
	RGB	2D	78.19±0.63	87.14±0.41
Linear	PCA	2D	82.06±1.01	89.34±0.65
	NMF	2D	82.93±1.63↑	89.93±1.17↑
	FactorAnalysis	2D	81.86±1.02	89.25±0.75
Non Linear	KPCA-RBF	2D	82.15±3.30	89.52±2.10
	t-SNE	2D	81.87±1.36	89.27±0.97
	Isomap	2D	83.28±0.87↑	90.24±0.62↑
	HSI	2D	77.96±2.11	87.05±1.31
		3D	73.28±1.05	83.92±0.80

Table 1 reports the quantitative segmentation performance obtained using different dimensionality reduction techniques. Overall, all dimensionality reduction methods improved overlapping chromosome segmentation compared with both the RGB and HSI baselines. While the RGB-based 2D U-Net achieved an IoU of 78.19% and a Dice score of 87.14%, all reduced hyperspectral representations yielded IoU values above 81% and Dice scores above 89%.

Among the evaluated linear methods, NMF achieved the highest performance, reaching an IoU of 82.93% and a Dice score of 89.93%. Within the nonlinear methods, Isomap obtained the best average performance, achieving the highest IoU 83.28% and Dice 90.24% across all evaluated approaches. In contrast, directly processing the complete hyperspectral cube using either a 2D or a 3D U-Net resulted in lower performance, indicating that reducing the spectral dimensionality before segmentation is more effective than directly exploiting the original hyperspectral data. These observations suggest that transforming the hyperspectral cube into a compact representation provides more discriminative features for overlapping chromosomes segmentation .

Figure 1 presents qualitative segmentation results obtained with the different input representations. The visual comparison confirms the quantitative trends reported in Table 1. The predictions obtained from the full HSI input, t-SNE, Kernel PCA, and Factor Analysis show less accurate segmentation masks, with

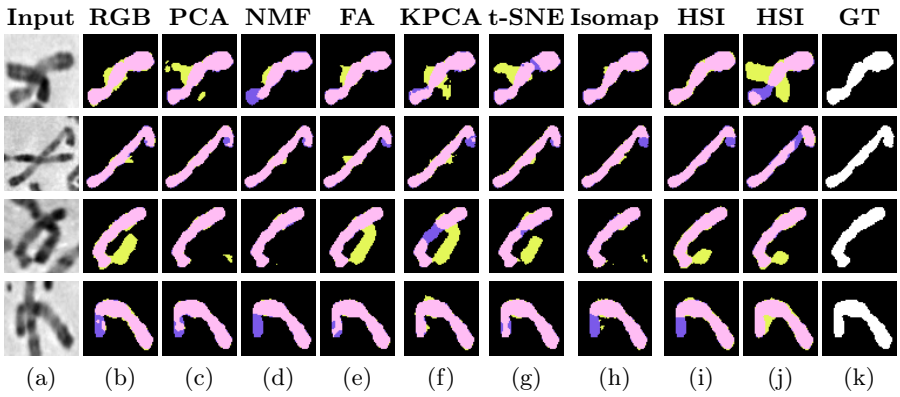

Fig. 1: Qualitative comparison of overlapping chromosome segmentation obtained using different dimensionality reduction techniques. The first column shows the input image, while the last column presents the corresponding ground truth. Pink-colored regions are correct predictions, purple-colored ones are false negatives, and yellow-colored ones are false positives. Panels (i) and (j) show the outputs of the 2D U-Net and 3D U-Net trained on the full HSI cube, respectively.

more visible false negatives and false positives around chromosome crossing regions.

3.2 Hyperparameter Sensitivity Analysis

Table 2 evaluates the sensitivity of the nonlinear dimensionality-reduction methods to their main hyperparameters. Kernel PCA exhibits the largest dependence on the kernel scale, with $\gamma_0/4$ achieving the highest performance. For t-SNE, a perplexity of 30 provides the highest average performance. Isomap remains comparatively stable across the investigated neighborhood sizes. Although $n_{\text{neighbors}} = 50$ yields the highest average scores. Overall, the hyperparameters affect the absolute scores and may slightly alter the average ranking, but the experiments do not provide evidence of a consistently superior nonlinear representation.

3.3 Statistical Analysis

To assess whether the observed improvements over the baseline inputs were statistically significant, paired Wilcoxon signed-rank tests were performed on the IoU scores. Figure 2 shows the IoU distributions using Violin plots, with the corresponding p-values annotated above each comparison.

The results indicate that all dimensionality reduction methods significantly outperform both the RGB baseline and the full HSI input with $p << 0.001$. This trend is observed for both linear and nonlinear methods, confirming that

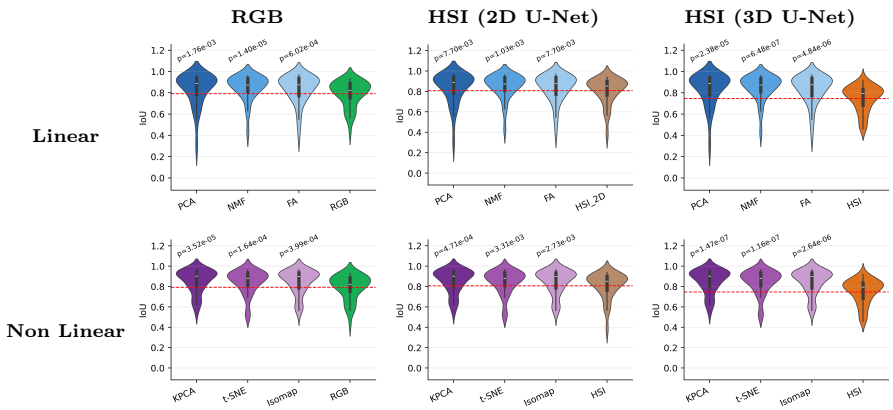

Fig. 2: Statistical analysis of IoU scores. Violin plots show the IoU distributions, and annotated values indicate the p-values obtained using the Wilcoxon signed-rank test. The first row compares the linear methods with the RGB, HSI baselines, while the second row presents the corresponding comparisons for the nonlinear methods.

Table 2: Sensitivity of nonlinear dimensionality-reduction methods to their main hyperparameters. $\gamma_0 = 1/104$. P. denotes perplexity and n. the number of neighbors.

KPCA-RBF			**t-SNE**			**Isomap**		
γ	**IoU**	**Dice**	**P.**	**IoU**	**Dice**	$n.$	**IoU**	**Dice**
$\gamma_0/4$	83.92±0.52 ↑	90.77±0.35 ↑	5	79.36±1.38	87.57±1.01	5	83.28±0.87	90.24±0.62
$\gamma_0/2$	82.42±1.91	89.70±1.29	10	80.30±0.45	88.22±0.25	10	82.93±0.59	89.88±0.43
γ_0	82.15±3.30	89.52±2.10	15	81.64±1.22	89.15±0.85	15	82.70±0.82	89.79±0.66
$2\gamma_0$	82.40±0.61	89.50±0.32	30	81.87±1.36↑	89.27±0.97↑	30	82.01±1.90	89.48±1.20
$4\gamma_0$	83.13±0.91	90.14±0.57	50	79.47±2.43	87.65±1.61	50	83.42±1.32 ↑	90.39±0.83 ↑

reducing the spectral dimensionality provides a consistent benefit for overlapping chromosomes segmentation. However, these statistical comparisons do not establish a clear superiority of linear over nonlinear dimensionality reduction methods. Instead, they show that both families of methods provide significant improvements over the baseline inputs. To further compare the dimensionality-reduction methods, pairwise Wilcoxon signed-rank tests were performed between all reduced representations. The resulting p-values were adjusted across the 15 comparisons using the Holm Bonferroni correction. As shown in Figure 3, none of the pairwise differences remains statistically significant after correction. Overall, the main finding of this study is that dimensionality reduction itself provides a consistent benefit for hyperspectral overlapping chromosome segmentation. By reducing the original hyperspectral cube to only three components, the proposed pipeline enables the use of a standard 2D U-Net while improving segmentation

performance compared with RGB and full HSI inputs. This is particularly relevant in biomedical applications, where annotated hyperspectral datasets are often limited and computational efficiency is important.

This study has some limitations. All dimensionality-reduction methods were restricted to three components to ensure a consistent comparison, and other dimensionalities may lead to different segmentation performance.

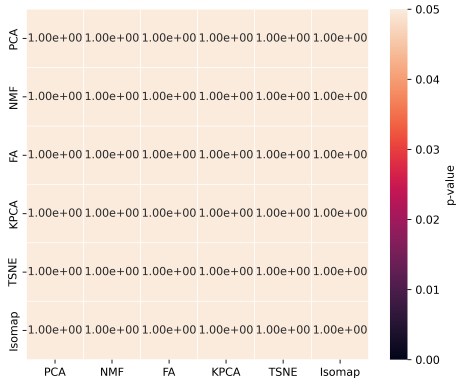

Fig. 3: Heatmap of p-values obtained from pairwise Wilcoxon signed-rank tests comparing the IoU scores of the evaluated dimensionality reduction methods. Lower p-values indicate stronger evidence of statistically significant differences between the corresponding method pairs.

4 Conclusion

This work investigated the impact of linear and nonlinear dimensionality reduction techniques on hyperspectral overlapping chromosome segmentation. Six representative methods were evaluated, including PCA, NMF, Factor Analysis, Kernel PCA, t-SNE, and Isomap, each producing a three-component representation used as input to a 2D U-Net. The results showed that dimensionality reduction significantly improves segmentation performance compared with both RGB images and the direct use of the full HSI cube with 2D and 3D U-Nets.

The main conclusion is that reducing the spectral dimensionality is beneficial for hyperspectral overlapping chromosome segmentation , while no clear statistical superiority of linear or nonlinear methods can be established from the current experiments.

Future work will investigate the influence of the number of retained components and evaluate the proposed framework on larger hyperspectral chromosome datasets.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)
2. Cunningham, J.P., Ghahramani, Z.: Linear dimensionality reduction: Survey, insights, and generalizations. The Journal of Machine Learning Research **16**(1), 2859–2900 (2015)
3. Fang, J.: A semi-supervised segmentation method for microscopic hyperspectral pathological images based on multi-consistency learning. Frontiers in Oncology **14**, 1396887 (2024)
4. Hotelling, H.: Analysis of a complex of statistical variables into principal components. Journal of educational psychology **24**(6), 417 (1933)
5. Huang, Z., Fujiwara, T., Chatzimparmpas, A., Duchemin, W., Kerren, A.: Maple: Self-supervised learning-enhanced nonlinear dimensionality reduction for visual analysis. arXiv preprint arXiv:2601.20173 (2026)
6. Lee, D.D., Seung, H.S.: Learning the parts of objects by non-negative matrix factorization. nature **401**(6755), 788–791 (1999)
7. Lu, G., Fei, B.: Medical hyperspectral imaging: a review. Journal of biomedical optics **19**(1), 010901–010901 (2014)
8. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008)
9. Rasti, B., Hong, D., Hang, R., Ghamisi, P., Kang, X., Chanussot, J., Benediktsson, J.A.: Feature extraction for hyperspectral imagery: The evolution from shallow to deep: Overview and toolbox. IEEE Geoscience and Remote Sensing Magazine **8**(4), 60–88 (2020)
10. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015)
11. Schölkopf, B., Smola, A., Müller, K.R.: Nonlinear component analysis as a kernel eigenvalue problem. Neural computation **10**(5), 1299–1319 (1998)
12. Seidlitz, S., Sellner, J., Odenthal, J., Özdemir, B., Studier-Fischer, A., Knödler, S., Ayala, L., Adler, T.J., Kenngott, H.G., Tizabi, M., et al.: Robust deep learning-based semantic organ segmentation in hyperspectral images. Medical Image Analysis **80**, 102488 (2022)
13. Spearman, C.: " general intelligence" objectively determined and measured. (1961)
14. Tenenbaum, J.B., Silva, V.d., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. science **290**(5500), 2319–2323 (2000)
15. Van Der Maaten, L., Postma, E.O., Van Den Herik, H.J., et al.: Dimensionality reduction: A comparative review. Journal of machine learning research **10**(1), 1–41 (2009)
16. Wang, R., Yao, M.: Overlapping chromosome segmentation model based on dual attention u-net network. In: Seventh International Conference on Advanced Electronic Materials, Computers, and Software Engineering (AEMCSE 2024). vol. 13229, pp. 416–421. SPIE (2024)
17. Wilcoxon, F.: Individual comparisons by ranking methods. In: Breakthroughs in statistics: Methodology and distribution, pp. 196–202. Springer (1992)

18. Yun, B., Lei, B., Chen, J., Wang, H., Qiu, S., Shen, W., Li, Q., Wang, Y.: Spectr: Spectral transformer for microscopic hyperspectral pathology image segmentation. IEEE Transactions on Circuits and Systems for Video Technology **34**(6), 4610–4624 (2023)
19. Yun, B., Li, Q., Mitrofanova, L., Zhou, C., Wang, Y.: Factor space and spectrum for medical hyperspectral image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 152–162. Springer (2023)
20. Zeffate, N., Bechar, M.E.A., Guyader, J.M., Settouti, N., Douet-Guilbert, N., Troadec, M.B., El Bouz, M.: Exploring band selection methods for enhanced chromosomal analysis in hyperspectral imaging. Applied Spectroscopy **80**(2), 133–145 (2026)
21. Zeffate, N., El Amine Bechar, M., Guyader, J.M., Douet-Guilbert, N., El Bouz, M., Troadec, M.B.: Exploring hsi imagery for overlapped chromosomes segmentation. Chromosome Research **33**(1), 30 (2025)
22. Zhou, T., Ruan, S., Canu, S.: A review: Deep learning for medical image segmentation using multi-modality fusion. Array **3**, 100004 (2019)