
Seeded iterative clustering for histology region identification

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Annotations are necessary to develop computer vision algorithms for histopathol-
2 ogy, but dense annotations at a high resolution are often time-consuming to make.
3 Deep learning models for segmentation are a way to alleviate the process, but
4 require large amounts of training data, training times and computing power. To
5 address these issues, we present *seeded iterative clustering* to produce a coarse
6 segmentation densely and at the whole slide level. The algorithm uses precomputed
7 representations as the clustering space and a limited amount of sparse interactive
8 annotations as seeds to iteratively classify image patches. We obtain a fast and
9 effective way of generating dense annotations for whole slide images and a frame-
10 work that allows the comparison of neural network latent representations in the
11 context of transfer learning.

12 1 Introduction

13 Deep learning segmentation models are already being used for easing the burden of cancer histopathol-
14 ogy whole slide image (WSI) annotation [22]. However, supervised training requires annotated
15 images, often long training times and computational resources such as graphical processing units
16 (GPU). To overcome these bottlenecks, approaches such as active learning [20] are gaining popularity,
17 as they reduce the time needed to perform the segmentation while retaining high performance via
18 expert supervision. One strategy is to have pathologists annotate a small part of the image, either by
19 clicking or scribbling around the regions of interest. Then, features are extracted from those sparsely
20 annotated regions and the dense segmentation masks are obtained by retraining segmentation models.
21 This strategy, however, requires both interactive and GPU support at run time.

22 For this proof-of-concept, we instead propose seeded interactive clustering (SIC) to decouple the
23 problem by (1) precomputing the latent network representations, which does not require a GPU at run
24 time so it can be done remotely in a computing cluster without interaction and, once the embeddings
25 are obtained, (2) perform interactive clustering locally that does not involve any further training and
26 is relatively cheaper computationally.

27 1.1 Previous work

28 Different authors have approached this problem by creating interactive tools based on extracting
29 classical features [1, 11] or on deep learning approaches [12, 14, 16]. Some of these approaches
30 are intended for cell detection on smaller images, and therefore are harder to scale to bigger WSI.
31 The approaches that are meant for digital pathology usually require a GPU while performing the
32 interactive annotations, which may not always be available.

33 The re-purposing of convolutional neural networks (CNN) as feature extractors was proposed very
 34 early after the first models trained on large datasets were made available [8, 18, 21, 23]. Since then,
 35 using model architectures such as ResNet-50 [10] trained on ImageNet has become common practice
 36 for extracting features and performing transfer learning, even if some authors argue that these models
 37 may not capture all the relevant histology features [4].

38 Recent work also proposes using self-supervised learning approaches for obtaining the embeddings [5,
 39 6], under the assumption that the feature space would be more representative of the tissue differences.
 40 However comparing different embeddings is not a trivial task and is influenced by many factors [13,
 41 17].

42 Finally, image-set clustering is a field where all these concepts, including the choice of architecture
 43 and layer for feature extraction, or the clustering algorithms, has been extensively studied [9].

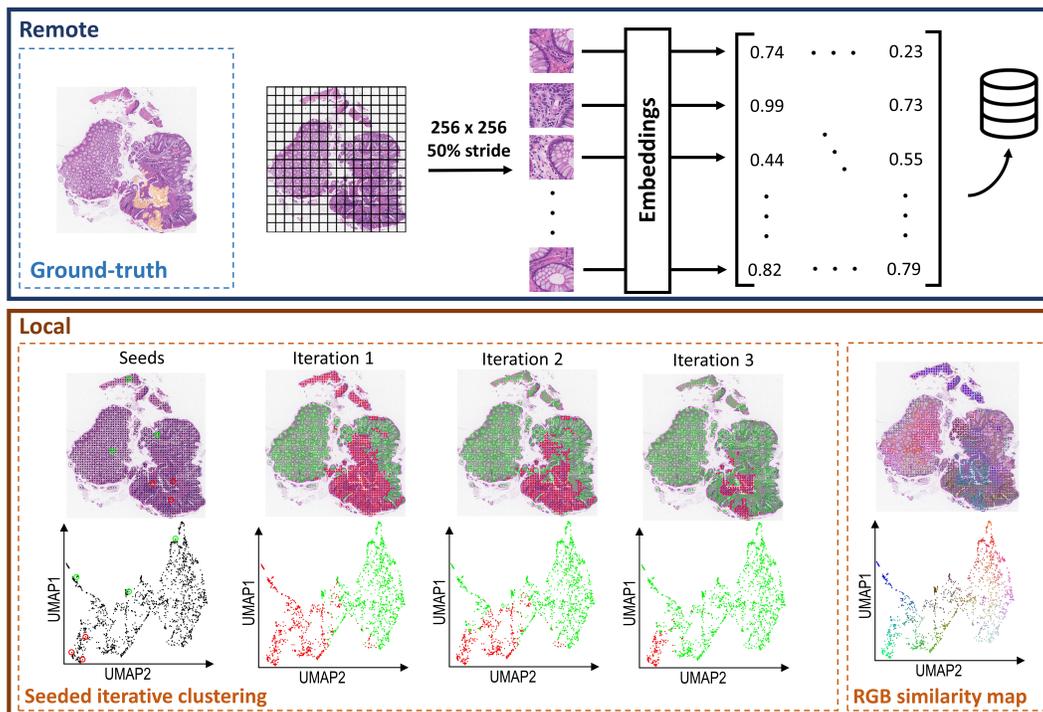


Figure 1: Visualization of the seeded interactive clustering process. The *Remote* box includes the ground truth (left) and illustrates how tissue patch embeddings are extracted (right). The *Local* box, from left to right: the top row shows a tissue sample, where seeds (red and green markers) are manually placed on the tissue image. The bottom row shows the embedding space location of the corresponding tissue patches, where black dots are patches not yet assigned to a category. As the SIC iterations continue, the patches are re-clustered according to embedding similarity. Far right: dots representing patches in UMAP space are colored according to their position in RGB color space, and re-mapped to their coordinates in tissue space. Note that the clustering is made on the full embedding space and the reduced space is only intended for visualization purposes. Visualizations created using TissUMaps [19].

44 2 Method

45 2.1 Seeded iterative clustering

46 The proposed pipeline is presented in Figure 1 and has two parts: one remote and one local. The
 47 first remote part of the process starts by dividing the image into smaller patches. The size of the

48 patches, the stride between them and the resolution of choice defines the granularity of the final tissue
 49 clustering. In this case, we decided to extract 256×256 patches with 50% overlap at $20\times$ resolution.
 50 Next, each of the patches are embedded using the pretrained model of choice by extracting the last
 51 layer latent representation. Finally, we store the matrix of *number of patches* \times *embedding size*
 52 together with the center coordinates of each patch to represent the complete WSI.

53 The second part of the pipeline can be run locally. The seeded iterative clustering (SIC) approach,
 54 presented in Algorithm 1, divides the previously generated patch embedding space in a binary fashion,
 55 guided by small seeds -or sparse annotations- selected by the user in representative regions. Then
 56 seeded clustering [2] is performed iteratively until a score relative to the sparse annotations does not
 57 improve. We used F1-score, but any classification metric could be used to optimize the algorithm.
 58 The seeded KMeans algorithm only differs from the original KMeans in the first assignment of cluster
 59 centers, which are not initialized randomly but based on the seeds. An additional step of conventional
 60 KMeans is added as it had shown to improve the performance when the tumor region of the WSI was
 61 not big enough.

Algorithm 1 Seeded Iterative Clustering

```

X  $\leftarrow$  embeddings                                ▷ Embeddings of patches
y  $\leftarrow$  seeds                                       ▷ Sparse annotations
scoreprev  $\leftarrow$  0                                     ▷ Initialize score
labels  $\leftarrow$  SeededKMeans(nclusters = 2).fit(X, y)    ▷ Initialize labels
while score  $\leq$  scoreprev do                               ▷ Run until defined score goes down
  X  $\leftarrow$  embeddings[labels == 1]                       ▷ Select data only in the positive class
  y  $\leftarrow$  seeds[labels == 1]
  if  $\{0, 1\} \subset y$  then                                       ▷ If there are still annotations of both classes
    labels  $\leftarrow$  SeededKMeans(nclusters = 2).fit(X, y)    ▷ Re-cluster positive cluster
  else
    if  $\{1\} \subset y$  then                                       ▷ If there are only negative annotations
      labels  $\leftarrow$  KMeans(nclusters = 2).fit(X)           ▷ Perform conventional clustering
    end if
  end if
  score  $\leftarrow$  metric(y, label)                          ▷ Compare only sparse annotations
end while

```

62 **2.2 Feature space visualization**

63 Inspired by the spatial-omics field, an intuitive way of visualizing the embeddings is by reducing
 64 the dimensionality of the space to three dimensions using Uniform Manifold Approximation and
 65 Projection for Dimension Reduction (UMAP) [15]. This 3D space can be mapped to the RGB space
 66 and visualized directly on top of the images as colored dots at the center of each patch in the image,
 67 as previously presented by Chelebian et al. [3] and shown in Figure 1 (bottom-right). The code for
 68 doing this can be found at github.com/eduardchelebian/histology-umap.

69 In the UMAP space, it is possible to manually select regions that are clustered together and see if they
 70 correspond to biologically relevant regions on the image. This approach however, is highly dependent
 71 on the dimensionality reduction algorithm and should be used only for qualitative visualization
 72 purposes.

73 **3 Experimental results**

74 **3.1 Dataset**

75 For this proof-of-concept we used the DigestPath challenge cancer segmentation dataset [7]. It is an
 76 open database including 250 images of tissue from 93 positive WSI (i.e. with pixel-level annotations)
 77 slides in JPEG format which were used for the experimental results. The additional 410 images
 78 from 231 WSI negative slides were only used together with the positive slides when self-supervised

79 pretraining was performed. All WSI are hematoxylin-eosin (H&E) stained and scanned at $20\times$
80 resolution.

81 3.2 Experiments

82 In order to explore how different embedding spaces would affect the clustering performance, we chose
83 four different feature extraction frameworks: (1) ResNet-18 pretrained on ImageNet (ResNet18),
84 (2) ResNet-50 pretrained on ImageNet (ResNet50), (3) ResNet-18 pretrained on the dataset using
85 SimCLR self-supervision [5] (SimCLR) and (4) ResNet-18 pretrained on 57 histopathology datasets
86 using SimCLR self-supervision [6] (HistoSSL).

87 To mimic the complete workflow where the pathologist would select small benign and cancerous
88 regions, we chose to experiment by randomly selecting 1, 3 and 5 patches that were located inside
89 the benign and cancerous masks, respectively. Knowing that this would probably underestimate the
90 results, as the seeds may not be representative of what a pathologist would choose, we report the
91 F1-score for the patch classification task of the best run per slide in a 5-fold experiment, to simulate
92 the selection an expert would make. The ground-truth coordinates are generated by selecting the
93 patches that are at least 50% in a cancer region.

94 3.3 Results

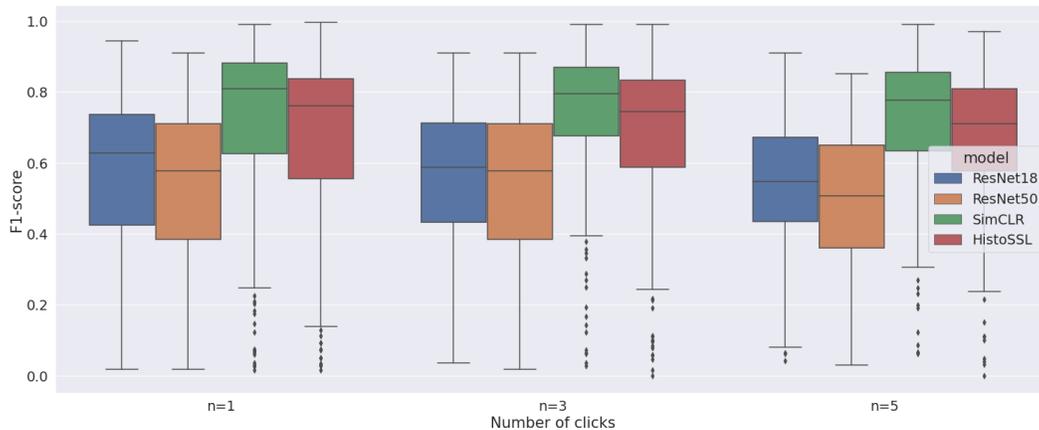


Figure 2: F1-score for the different models with varying number of simulated clicked patches.

95 The results are summarized in Figure 2. The highest per-slide F1-score for the ImageNet pretrained
96 models are very similar regardless of the architecture, 0.57 ± 0.22 for ResNet18 and 0.54 ± 0.21
97 for ResNet50, both lower than the self-supervised approaches. There seems to be some benefit for
98 pretraining (SimCLR) on the dataset at hand (0.74 ± 0.23), but this may not be always possible, so it is
99 very promising that a publicly available model (HistoSSL) also achieves higher results (0.68 ± 0.22).
100 The method itself shows high classification variability, although the worse performing slides are
101 usually the ones that have low positive patch ratio, that is, when the lesion is very small compared to
102 the whole tissue, such as Figure 3 (upper right). There is no clear difference from having more seeds,
103 except when it comes to the standard deviation. As they are only used for initialization it seems
104 that if the feature space is descriptive enough, adequately selected patches are sufficient for a good
105 clustering result.

106 4 Discussion

107 We presented seeded iterative clustering (SIC) as a way of leveraging latent representations of neural
108 networks to speed-up the time-consuming process of manual annotation in histopathology. From an

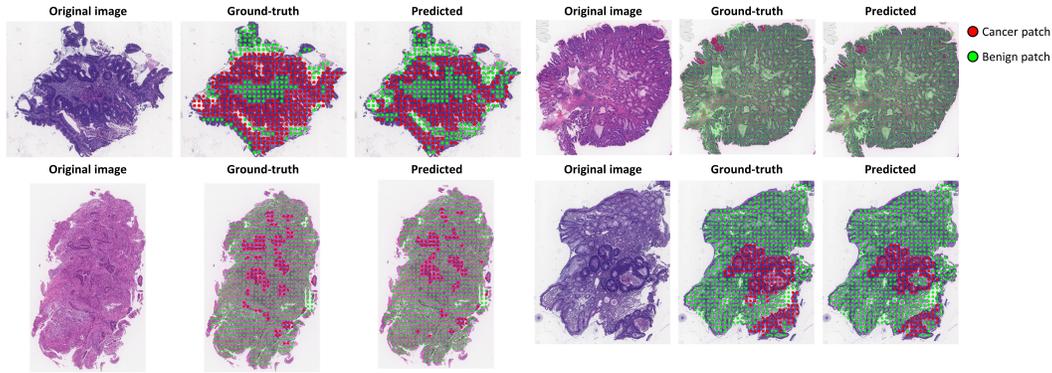


Figure 3: Examples of predictions using seeded iterative clustering. Visualizations created using TissUUmeps [19].

109 annotation tool point of view, the local clustering module takes seconds, even on a normal computer.
 110 The embedding time depends on the amount of patches or coordinates one wants. However, contrary
 111 to previous work, this can be run remotely without the pathologist interaction and does not need to be
 112 recalculated at any point. Additionally, performing the clustering per-slide is more robust to domain
 113 shift artifacts.

114 From a technical point of view, it allowed to compare the representations from different networks.
 115 When performing transfer learning, there is a separate effect of the feature reuse and learning low-level
 116 statistics from the data [17]. Retraining, then, may not be ideal to explore the quality of the features.
 117 This is precisely why we find that ImageNet pretrained ResNet off-the-shelf are not comparable to
 118 using self-supervised pretraining with SimCLR.

119 Work in progress includes validating this approach on other datasets and developing it into a full
 120 interactive tool called in which pathologists select a handful of patches and the method proposes
 121 clusters for the rest. Additionally, the method needs to be compared against other approaches in
 122 order to confirm its feasibility. Finally, extending this approach in a multi-class problem (e.g. cancer,
 123 benign and stroma regions, instead of binary), would provide further insights into how the feature
 124 space generated by the different models as well as increase the usability of the tool.

125 **Broader Impact**

126 The uncoupling of the feature extraction and clustering modules could allow for the remote generation
 127 of embeddings that can then be shared with the community. This would allow for groups with less
 128 computational resources to work with the data and circumvent the potential privacy issues for sharing
 129 medical images, sharing instead the latent representation matrices per coordinate.

130 **References**

- 131 [1] Arzt, M., Deschamps, J., Schmied, C., Pietzsch, T., Schmidt, D., Tomancak, P., Haase, R.,
 132 & Jug, F. (2022). Labkit: Labeling and segmentation toolkit for big image data. *Frontiers in*
 133 *Computer Science*, 10.
 134 [2] Basu, S., Banerjee, A., & Mooney, R. (2002). Semi-supervised clustering by seeding. *In*
 135 *Proceedings of 19th International Conference on Machine Learning (ICML-2002)*.
 136 [3] Chelebian, E., Avenel, C., Kartasalo, K., Marklund, M., Tanoglidi, A., Mirtti, T., Colling, R.,
 137 Erickson, A., Lamb, A. D., Lundeberg, J., et al. (2021). Morphological features extracted by ai
 138 associated with spatial transcriptomics in prostate cancer. *Cancers*, 13(19), 4837.
 139 [4] Chen, R. J., & Krishnan, R. G. (2022). Self-supervised vision transformers learn visual
 140 concepts in histopathology. *arXiv preprint arXiv:2203.00585*.

- 141 [5] Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive
142 learning of visual representations. *International conference on machine learning*, 1597–1607.
- 143 [6] Ciga, O., Xu, T., & Martel, A. L. (2022). Self supervised contrastive learning for digital
144 histopathology. *Machine Learning with Applications*, 7, 100198.
- 145 [7] Da, Q., Huang, X., Li, Z., Zuo, Y., Zhang, C., Liu, J., Chen, W., Li, J., Xu, D., Hu, Z., et al.
146 (2022). Digestpath: A benchmark dataset with challenge review for the pathological detection
147 and segmentation of digestive-system. *Medical Image Analysis*, 102485.
- 148 [8] Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., & Darrell, T. (2014).
149 Decaf: A deep convolutional activation feature for generic visual recognition. *International
150 conference on machine learning*, 647–655.
- 151 [9] Guérin, J., Thiery, S., Nyiri, E., Gibaru, O., & Boots, B. (2021). Combining pretrained cnn
152 feature extractors to enhance clustering of complex natural images. *Neurocomputing*, 423,
153 551–571.
- 154 [10] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition.
155 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- 156 [11] Hollandi, R., Diósdí, Á., Hollandi, G., Moshkov, N., & Horváth, P. (2020). Annotatorj: An
157 imagej plugin to ease hand annotation of cellular compartments. *Molecular biology of the cell*,
158 31(20), 2179–2186.
- 159 [12] Jaber, M. I., Song, B., Beziaeva, L., Szeto, C. W., Spilman, P., Yang, P., & Soon-Shiong, P.
160 (2021). A deep learning-based iterative digital pathology annotation tool. *bioRxiv*.
- 161 [13] Kornblith, S., Norouzi, M., Lee, H., & Hinton, G. (2019). Similarity of neural network
162 representations revisited. *International Conference on Machine Learning*, 3519–3529.
- 163 [14] Lindvall, M., Sanner, A., Petré, F., Lindman, K., Treanor, D., Lundstrbm, C., & Ldwgren, J.
164 (2020). Tissuewand, a rapid histopathology annotation tool. *Journal of pathology informatics*,
165 11(1), 27.
- 166 [15] McInnes, L., Healy, J., & Melville, J. (2018). Umap: Uniform manifold approximation and
167 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- 168 [16] Miao, R., Toth, R., Zhou, Y., Madabhushi, A., & Janowczyk, A. (2021). Quick annotator: An
169 open-source digital pathology based rapid image annotation tool. *The Journal of Pathology:
170 Clinical Research*, 7(6), 542–547.
- 171 [17] Neyshabur, B., Sedghi, H., & Zhang, C. (2020). What is being transferred in transfer learning?
172 *Advances in neural information processing systems*, 33, 512–523.
- 173 [18] Oquab, M., Bottou, L., Laptev, I., & Sivic, J. (2014). Learning and transferring mid-level image
174 representations using convolutional neural networks. *Proceedings of the IEEE conference on
175 computer vision and pattern recognition*, 1717–1724.
- 176 [19] Pielawski, N., Andersson, A., Avenel, C., Behanova, A., Chelebian, E., Klemm, A., Nysjö,
177 F., Solorzano, L., & Wählby, C. (2022). Tissuumaps 3: Interactive visualization and quality
178 assessment of large-scale spatial omics data. *bioRxiv*.
- 179 [20] Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., Chen, X., & Wang, X. (2021).
180 A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9), 1–40.
- 181 [21] Sharif Razavian, A., Azizpour, H., Sullivan, J., & Carlsson, S. (2014). Cnn features off-the-
182 shelf: An astounding baseline for recognition. *Proceedings of the IEEE conference on computer
183 vision and pattern recognition workshops*, 806–813.
- 184 [22] Van der Laak, J., Litjens, G., & Ciompi, F. (2021). Deep learning in histopathology: The path
185 to the clinic. *Nature medicine*, 27(5), 775–784.
- 186 [23] Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks.
187 *European conference on computer vision*, 818–833.