refinement of lesion detection metrics in the presence of confluent lesions

Author: BenchChem Technical Support Team. **Date**: December 2025

Compound of Interest		
Compound Name:	WYneN	
Cat. No.:	B15548174	Get Quote

Technical Support Center: Lesion Detection Metrics

Welcome to the technical support center for lesion detection and segmentation analysis. This resource provides troubleshooting guides and answers to frequently asked questions (FAQs) regarding the evaluation of models, particularly in challenging scenarios involving confluent lesions.

Frequently Asked questions (FAQs)

Q1: What are confluent lesions and why do they pose a challenge for automated detection algorithms?

A: Confluent lesions are distinct pathological abnormalities that have grown close together and merged, appearing as a single, larger lesion in medical images.[1][2] This confluence is a significant challenge for automated analysis because it complicates the fundamental tasks of counting and individually segmenting lesions.[2][3] For algorithms, what is clinically understood as multiple distinct lesions may be computationally seen as a single object, leading to inaccuracies in both quantitative and qualitative assessments.[2]

Q2: How do standard, pixel-based metrics like Intersection over Union (IoU) and Dice Coefficient fail when evaluating confluent lesions?

A: Standard metrics like Intersection over Union (IoU) and the Dice Coefficient are pixel-overlap-based; they measure how well a predicted segmentation mask matches the ground truth mask. While excellent for assessing the accuracy of a single object's boundary, they are fundamentally flawed for evaluating confluent lesions.

The primary issue is that these metrics cannot distinguish between the detection of one large lesion versus several smaller, individual lesions that are correctly identified. For example, if two distinct lesions in the ground truth are detected as one single merged lesion by the model, the IoU or Dice score can still be very high because the overall pixel overlap is substantial. This masks the model's failure to perform correct instance segmentation (identifying each lesion as a separate object).

Q3: What are more refined, object-level metrics that are better suited for this challenge?

A: To properly evaluate performance in the presence of confluent lesions, it is crucial to use object-level (or instance-level) metrics that assess both detection and segmentation quality. A highly effective paradigm for this is the Free-Response Receiver Operating Characteristic (FROC) analysis.

- FROC Analysis: Unlike standard ROC curves that provide an image-level classification (lesion present/absent), FROC evaluates the model's ability to correctly localize each lesion. It plots the Lesion Localization Fraction (LLF) against the Non-Lesion Localization Fraction (NLF) or the rate of false positives per image. This approach rewards the correct identification of each individual lesion, making it far more insightful for confluent scenarios.

- Jackknife Alternative FROC (JAFROC): This is an advanced method for analyzing FROC data that has been shown to have greater statistical power than standard ROC analysis for tasks involving lesion localization.

- Instance Segmentation Metrics: Metrics adapted from computer vision tasks, such as Panoptic Quality (PQ) and Mean Average Precision (mAP) at various IoU thresholds, are also valuable. These metrics evaluate a model's ability to correctly detect and segment each individual object instance.

Q4: What is the difference between semantic and instance segmentation in the context of lesion analysis?

A:

- Semantic Segmentation: This process assigns a class label (e.g., "lesion" or "not lesion") to every pixel in an image. It produces a single mask for all lesions and cannot distinguish between individual, touching lesions.

- Instance Segmentation: This is a more advanced task that not only classifies each pixel but also identifies which object instance each pixel belongs to. In the context of confluent lesions, an ideal instance segmentation model would produce separate, distinct masks for each individual lesion, even if they are touching.

Troubleshooting Guides

Problem 1: My algorithm detects multiple confluent lesions as a single object. How do I score this correctly?

Symptom: Your model produces one large bounding box or segmentation mask that covers two or more adjacent lesions defined in the ground truth. A simple IoU or Dice calculation returns a high score, suggesting good performance, but visual inspection shows a failure to separate the instances.

Troubleshooting Steps:

- Shift from Pixel-Overlap to Object-Level Evaluation: Do not rely solely on Dice or IoU. These metrics are misleading here. Implement an FROC analysis workflow. This requires matching each predicted lesion to a ground truth lesion. A prediction is typically counted as a "lesion localization" or true positive if its center falls within a certain radius of a ground truth lesion's center or if the IoU between the prediction and a ground truth lesion exceeds a set threshold (e.g., 0.5).

- Penalize the Missed Detections: In the scenario described, the single large prediction would be matched to one of the ground truth lesions (counting as one true positive), while the other ground truth lesion(s) it covers would be counted as false negatives. This will be accurately reflected in object-level metrics like precision, recall, and the F1-score for detection.

- Visualize the Discrepancy: The diagram below illustrates how a high IoU can mask a detection failure.

Caption: The Confluence Problem: High pixel overlap (IoU) can mask poor instance detection.

Problem 2: How should I annotate ground truth data when lesions are touching or overlapping?

Symptom: You are unsure whether to label a confluent region as one large mask or as multiple individual-but-overlapping masks. This decision critically impacts how any algorithm will be trained and evaluated.

Troubleshooting Steps & Protocol:

- Adopt an Instance-Based Annotation Protocol: The best practice is to annotate each clinically distinct lesion as a separate instance, even if the pixels overlap. This creates the necessary ground truth for training and validating instance segmentation models.

- Use Appropriate Tooling: Employ annotation software that is designed for medical imaging and supports overlapping polygon or mask annotations on different layers or with distinct instance IDs. General-purpose tools may not handle this correctly.

- Establish Clear Guidelines with Clinical Experts: Work directly with radiologists or other domain experts to define what constitutes a separable lesion versus a single, large lesion. Document these rules with visual examples of edge cases to ensure consistency across all annotators.

- Protocol: Annotating Confluent Lesions:

  - Step 1: Initial Identification: A clinical expert identifies a region of confluent lesions on a scan (e.g., MRI, CT).

  - Step 2: Instance-by-Instance Delineation: Using a tool that supports multiple layers or instance IDs, the annotator carefully draws the boundary for the first lesion.

  - Step 3: Overlapping Annotation: The annotator then draws the boundary for the second lesion, allowing the masks to overlap where necessary. Each mask is saved with a unique ID.

  - Step 4: Expert Review: A second expert should review a subset of annotations to ensure adherence to the established guidelines and maintain high inter-annotator agreement.

  - Step 5: Finalization: The final ground truth is saved in a format that preserves the individual instance masks (e.g., COCO format for instance segmentation).

Problem 3: Which metric should I choose for my study? The choice seems to impact my results significantly.

Symptom: You have calculated multiple metrics (e.g., Dice, FROC, mAP) and they tell different stories about your model's performance. You need to decide which is most relevant for your research or clinical application.

Troubleshooting Steps:

- Align Metric with Clinical Goal: The choice of metric should be driven by the clinical question you are trying to answer.

  - For tracking lesion count/burden (e.g., in Multiple Sclerosis): Object-level detection metrics are critical. FROC or a detection F1-score will tell you if your model is accurately counting new or changing lesions.

  - For assessing treatment response based on lesion volume: A combination is needed. You need a good detection metric to ensure you're measuring the right lesions, and a good segmentation metric (like Dice) to ensure the volume calculation is accurate for the correctly identified lesions.

  - For screening and diagnosis: Sensitivity at a low false-positive rate is paramount. FROC analysis is the standard here as it directly visualizes this trade-off.

- Use a Combination of Metrics: A single metric is rarely sufficient. Best practice is to report a holistic view:

  - An object-level detection metric (e.g., FROC, JAFROC, mAP).

  - A pixel-level segmentation metric for the correctly detected lesions (e.g., Dice or IoU).

- Decision Workflow: The following diagram provides a logical workflow for selecting the appropriate evaluation strategy.

Caption: Logic diagram for selecting the appropriate lesion evaluation metric based on study goals.

Data Summary: Metric Comparison

The following table summarizes the key characteristics of common metrics and their suitability for evaluating confluent lesions.

Metric	Type	What It Measures	Suitability for Confluent Lesions
Intersection over Union (IoU) / Jaccard	Pixel-based Overlap	The ratio of the intersection to the union of predicted and ground truth masks.	Poor. Can be high even when individual lesions are incorrectly merged, masking detection failures.
Dice Coefficient / F1 Score	Pixel-based Overlap	The harmonic mean of precision and recall at the pixel level; highly correlated with IoU.	Poor. Suffers from the same limitations as IoU regarding merged detections.
Free-Response ROC (FROC)	Object-based Detection	Plots lesion localization fraction vs. false positives per image, assessing detection and localization.	Excellent. Directly evaluates the model's ability to identify each lesion instance separately.
Mean Average Precision (mAP)	Object-based Detection & Segmentation	Averages precision across recall values and IoU thresholds for each object instance.	Very Good. Standard for instance segmentation; effectively penalizes both missed and poorly segmented lesions.
Panoptic Quality (PQ)	Hybrid (Detection & Segmentation)	Combines segmentation quality (IoU) and detection quality (F1-score) into a single metric.	Very Good. Provides a unified score that explicitly accounts for both instance detection and segmentation accuracy.

> **Need Custom Synthesis?**
>
> *BenchChem offers custom synthesis for rare earth carbides and specific isotopiclabeling.*
>
> *Email: info@benchchem.com or Request Quote Online.*

References

- 1. brainrehabilitation.org [brainrehabilitation.org]
- 2. An Automated Statistical Technique for Counting Distinct Multiple Sclerosis Lesions - PMC [pmc.ncbi.nlm.nih.gov]
- 3. Conflunet: Improving Confluent Lesion Identification In Multiple Sclerosis With Instance Segmentation | IEEE Conference Publication | IEEE Xplore [ieeexplore.ieee.org]
- To cite this document: BenchChem. [refinement of lesion detection metrics in the presence of confluent lesions]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15548174#refinement-of-lesion-detection-metrics-in-the-presence-of-confluent-lesions]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [Contact our Ph.D. Support Team for a compatibility check]

Need Industrial/Bulk Grade?   Request Custom Synthesis Quote

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd

Ontario, CA 91761, United States

Phone: (601) 213-4426

Email: info@benchchem.com

Tech Support