
Breast Tumor Segmentation in Dynamic Contrast-Enhanced MRI via Multi-Staged Training and Deep Ensembling of a Large Kernel MedNeXt

Anonymous Author(s)

Affiliation

Address

email

1 1 Introduction

2 Breast cancer is a major health concern [8], and Magnetic Resonance Imaging (MRI) plays an
3 important role in its assessment, preoperative staging and treatment [14, 7]. T1-weighted dynamic
4 contrast-enhanced (DCE) MRI can highlight tumor vascularity by using contrast agents, aiding
5 in the localization of tumor lesions. Accurate segmentation of tumor boundaries on these images
6 is clinically valuable, as it enables quantitative evaluation of tumor size, shape, and volume over
7 time [16, 1]. High-quality tumor segmentation further facilitates advanced analyses such radiomics
8 feature extraction [17] for other downstream tasks [2, 4], including pCR assessment [13] and the
9 characterization of tumor types [5].

10 Manual biomedical image segmentation is highly time-consuming, and tends to suffer from inter-
11 and intra-annotator variability subject to the level of experience of a radiologist. Deep learning
12 based automated segmentation methods have proven to be reliable in addressing the stated problems
13 [9, 15, 12, 11], with some limitations in terms of data, and or architectural constraints. To achieve
14 robust and generalizable performances, highly parameterized deep learning models need to be trained
15 with sufficient enough labeled data. Readily available large-scale labeled medical imaging data is
16 lacking, particularly for dynamic contrast-enhanced MRI segmentation.

17 We address the challenge of the segmentation of breast tumor lesions in multi-contrast MRI, using
18 MedNeXt and a multi-staged training strategy of improving receptive field, and loss optimization
19 based deep ensembling. Our segmentation method is motivated by the need to capture both the
20 subtle enhancement patterns of tumors across multiple post-contrast time points and the broader
21 breast tissue context, which larger receptive fields naturally accommodate. To effectively utilize large
22 kernels without overfitting, we adopt a two-stage training strategy: we first train a MedNeXt model
23 with the conventional $3 \times 3 \times 3$ kernel sizes, and then expand to $5 \times 5 \times 5$ kernel sizes via trilinear
24 interpolation [15].

25 2 Methodology

26 2.1 Dataset and Preprocessing

27 We used a multicenter dataset of 1506 cases from over 20 institutions for training, and a held out 58
28 cases for testing [6]. Each case in the dataset includes a series of T1-weighted DCE-MRI volumes
29 acquired at multiple time points: one pre-contrast and up to five post-contrast phases. For our
30 experiments, we selected the pre-contrast image and the first two post-contrast images of each case.

31 The dataset comprises both unilateral and bilateral breast DCE-MRI scans. Data preprocessing,
32 training and inference were done using the standard nnU-Net pipeline [9]. During training, we
33 adopted nnU-Net’s patch-based sampling strategy with a fixed input size of $128 \times 128 \times 128$ voxels.

34 To ensure sufficient exposure to tumor regions across both scan types, more than one-third of
 35 the patches in each batch were enforced to contain at least one foreground voxel. At inference
 36 time, segmentation was performed using overlapping sliding-window patches, ensuring full-volume
 37 coverage regardless of laterality.

38 **2.2 MedNeXt Architecture**

39 MedNeXt is a fully ConvNeXt [10] encoder-decoder U-shaped network for biomedical image
 40 segmentation [15]. Its inverted bottleneck design in the up and downsampling layers, and the
 41 compound scaling of depth, width and receptive field, makes it a highly capable segmentation
 42 method. It is further transformer-inspired in its scaling approach, and the use of large-kernel sizes
 43 to approximate attention. The added inductive bias provides the benefits of both convolutional- and
 44 transformer-based approaches, in capturing short and long range dependencies respectively.

$$M^{(5)} = UpKern(M^{(3)}, \text{size} = 5) \quad (1)$$

45 The approximation of attention via larger kernel sizes of $5 \times 5 \times 5$ instead of the conventional $3 \times 3 \times 3$
 46 sizes is achieved by first pretraining a conventional MedNeXt (M^3), and trilinearly interpolating
 47 its convolutional kernels to an initialized large-kernel MedNeXt (M^5), using an algorithm called
 48 *UpKern* [15] in Equation 1. The performance saturation usually observed with increasingly large
 49 kernel sizes is mitigated [3].

50 **2.3 Training Strategy**

51 All networks were trained with deep supervision, stochastic gradient descent (SGD) optimization,
 52 and a cosine annealing learning rate schedule initialized at 1×10^{-4} , using an A100 NVIDIA GPU.
 53 For the base M^3_{Base} , training was conducted using a five-fold cross-validation split and optimized
 54 using Dice cross-entropy loss. Each fold was trained independently for 250 epochs. Following the
 55 completion of cross-validation, we identified the fold that achieved the highest mean Dice coefficient
 56 on its respective validation set, whose weights were used in the second stage of training.
 57 In the subsequent stage, we employed the *UpKern* strategy 1 to resample the learned M^3_{Base}
 58 convolutional kernel weights into M^5_{Base} via trilinear interpolation. This approach enabled a smooth
 59 transition to a large-kernel configuration, thus expanding the effective receptive field of the network
 60 without introducing instability often associated with training large kernels from random initialization
 61 [3]. The newly initialized M^5_{Base} was then fine-tuned for making use of the entire training set, with
 62 all other architectural and training settings held constant.
 63 Additionally, we generated a second M^5 by applying the *UpKern* algorithm and fine-tuning the
 64 pretrained weights, forming M^5_{Focal} . In this stage, the network was optimized using a composite
 65 loss that combines Dice–cross-entropy and focal loss to better penalize small lesion segmentation
 66 errors and address class imbalance arising from large foreground–background differences:

$$\mathcal{L}_{\text{total}} = 0.25 \mathcal{L}_{\text{Dice-CE}} + 0.75 \mathcal{L}_{\text{Focal}} \quad (2)$$

67 where $\mathcal{L}_{\text{Dice-CE}}$ denotes the combined Dice and cross-entropy loss used in the first training stage, and
 68 $\mathcal{L}_{\text{Focal}}$ is the focal loss component that increases the weighting of hard-to-segment regions. Code is
 69 publicly available along with the implemented composite loss functions¹.

70 **3 Results**

71 Segmentation performance was evaluated on the held-out testing set. The results are summarized
 72 in Table 1. M^3 achieved a Dice score of 0.64 and a normalized Hausdorff Distance (NormHD) of
 73 0.3. Upon applying the *UpKern* strategy and fine-tuning the best performing single model, M^5_{Base}
 74 improved the Dice score to 0.66 and the NormHD to 0.29 (see Figure 1). Further ensembling with
 75 M^5_{Focal} led to minimal increase in Dice to 0.67, and NormHD to 0.24. The reported baseline is a
 76 5-fold nnU-Net ensemble by [6], also trained on the 1506 training dataset and evaluated on the test
 77 set.

¹<https://github.com/toufiqmusah/caladan-mama-mia>

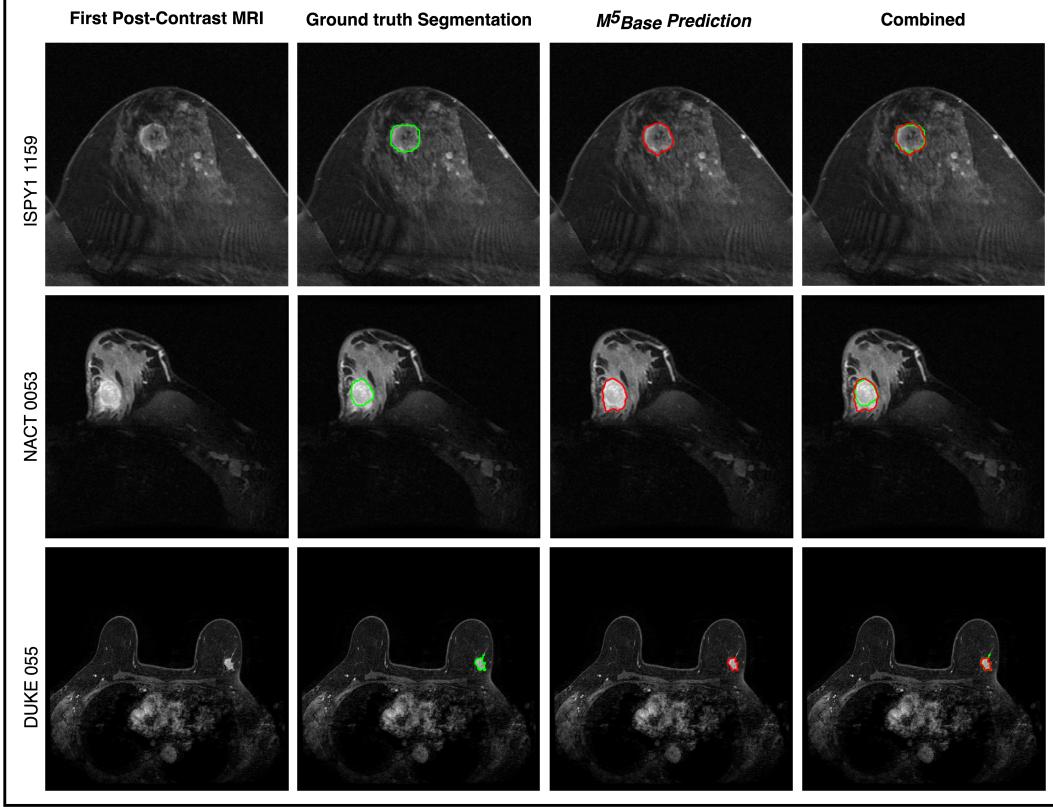


Figure 1: Qualitative segmentation performance of M^5_{Base} on validation samples from various centers in red, and ground truth in green.

Table 1: Performance comparison between methods. Metrics are reported as Dice Score and 95th-Percentile Normalized Hausdorff Distance. Total number of combined parameters are denoted in millions.

Architecture	Dice Score (\uparrow)	NormHD (\downarrow)	Parameters (\downarrow) (M)
nnU-Net (5-Folds) [6]	0.65	0.30	700
M^3 (5-Folds)	0.64	0.30	154
M^5_{Base}	0.66	0.29	32.1
$M^5_{Base} + M^5_{Focal}$	0.67	0.24	64.2

78 The 5-fold M^3 ensemble’s performance improved after applying the *UpKern* strategy to expand the
 79 receptive field. Fine-tuning the single best-performing fold into M^5_{Base} led to a Dice score of 0.66
 80 and a reduction in NormHD to 0.29. Further ensembling M^5_{Base} with M^5_{Focal} , which was trained
 81 using Dice-cross-entropy and focal loss, resulted in a Dice score of 0.67 and a NormHD of 0.24.
 82 These results validate the hypothesis that larger receptive fields improve segmentation performance
 83 by capturing broader anatomical context. Notably, both M^5_{Base} and $M^5_{Base} + M^5_{Focal}$ outper-
 84 formed the nnU-Net baseline [6] and achieved a lower NormHD, despite having substantially fewer
 85 parameters per model (32.1M vs. 140M per model instance) and using only two models in the final
 86 ensemble compared to the five-model nnU-Net baseline. Architectural efficiency and targeted loss
 87 function strategies can deliver improved performance while reducing computational requirements.
 88 Our study is limited in scope to the evaluation of the proposed *UpKern* strategy within the MedNeXt
 89 architecture; we did not test its generalizability across other network families. While ensemble
 90 performance was reported, we did not isolate and report the standalone performance of individual
 91 models within the ensemble, which could provide further insights into complementarity.

92 **NeurIPS Paper Checklist**

93 **1. Claims**

94 Question: Do the main claims made in the abstract and introduction accurately reflect the
95 paper's contributions and scope?

96 Answer: **[Yes]**

97 Justification: The introduction summarizes the contributions of a two stage training strategy
98 and ensembling of large-kernel MedNeXt.

99 Guidelines:

100 • The answer NA means that the abstract and introduction do not include the claims
101 made in the paper.

102 • The abstract and/or introduction should clearly state the claims made, including the
103 contributions made in the paper and important assumptions and limitations. A No or
104 NA answer to this question will not be perceived well by the reviewers.

105 • The claims made should match theoretical and experimental results, and reflect how
106 much the results can be expected to generalize to other settings.

107 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
108 are not attained by the paper.

109 **2. Limitations**

110 Question: Does the paper discuss the limitations of the work performed by the authors?

111 Answer: **[Yes]**

112 Justification: The paper discusses the limitations of not reporting on individual ensemble
113 component gains.

114 Guidelines:

115 • The answer NA means that the paper has no limitation while the answer No means that
116 the paper has limitations, but those are not discussed in the paper.

117 • The authors are encouraged to create a separate "Limitations" section in their paper.

118 • The paper should point out any strong assumptions and how robust the results are to
119 violations of these assumptions (e.g., independence assumptions, noiseless settings,
120 model well-specification, asymptotic approximations only holding locally). The authors
121 should reflect on how these assumptions might be violated in practice and what the
122 implications would be.

123 • The authors should reflect on the scope of the claims made, e.g., if the approach was
124 only tested on a few datasets or with a few runs. In general, empirical results often
125 depend on implicit assumptions, which should be articulated.

126 • The authors should reflect on the factors that influence the performance of the approach.
127 For example, a facial recognition algorithm may perform poorly when image resolution
128 is low or images are taken in low lighting. Or a speech-to-text system might not be
129 used reliably to provide closed captions for online lectures because it fails to handle
130 technical jargon.

131 • The authors should discuss the computational efficiency of the proposed algorithms
132 and how they scale with dataset size.

133 • If applicable, the authors should discuss possible limitations of their approach to
134 address problems of privacy and fairness.

135 • While the authors might fear that complete honesty about limitations might be used by
136 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
137 limitations that aren't acknowledged in the paper. The authors should use their best
138 judgment and recognize that individual actions in favor of transparency play an impor-
139 tant role in developing norms that preserve the integrity of the community. Reviewers
140 will be specifically instructed to not penalize honesty concerning limitations.

141 **3. Theory assumptions and proofs**

142 Question: For each theoretical result, does the paper provide the full set of assumptions and
143 a complete (and correct) proof?

144 Answer: [Yes]

145 Justification: The paper demonstrates increased segmentation performance with UpKern
146 and loss-optimized ensembles.

147 Guidelines:

- 148 • The answer NA means that the paper does not include theoretical results.
- 149 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
150 referenced.
- 151 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 152 • The proofs can either appear in the main paper or the supplemental material, but if
153 they appear in the supplemental material, the authors are encouraged to provide a short
154 proof sketch to provide intuition.
- 155 • Inversely, any informal proof provided in the core of the paper should be complemented
156 by formal proofs provided in appendix or supplemental material.
- 157 • Theorems and Lemmas that the proof relies upon should be properly referenced.

158 4. Experimental result reproducibility

159 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
160 perimental results of the paper to the extent that it affects the main claims and/or conclusions
161 of the paper (regardless of whether the code and data are provided or not)?

162 Answer: [Yes]

163 Justification: All information needed for reproducibility is provided in the methodology.

164 Guidelines:

- 165 • The answer NA means that the paper does not include experiments.
- 166 • If the paper includes experiments, a No answer to this question will not be perceived
167 well by the reviewers: Making the paper reproducible is important, regardless of
168 whether the code and data are provided or not.
- 169 • If the contribution is a dataset and/or model, the authors should describe the steps taken
170 to make their results reproducible or verifiable.
- 171 • Depending on the contribution, reproducibility can be accomplished in various ways.
172 For example, if the contribution is a novel architecture, describing the architecture fully
173 might suffice, or if the contribution is a specific model and empirical evaluation, it may
174 be necessary to either make it possible for others to replicate the model with the same
175 dataset, or provide access to the model. In general, releasing code and data is often
176 one good way to accomplish this, but reproducibility can also be provided via detailed
177 instructions for how to replicate the results, access to a hosted model (e.g., in the case
178 of a large language model), releasing of a model checkpoint, or other means that are
179 appropriate to the research performed.
- 180 • While NeurIPS does not require releasing code, the conference does require all submis-
181 sions to provide some reasonable avenue for reproducibility, which may depend on the
182 nature of the contribution. For example
 - 183 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
184 to reproduce that algorithm.
 - 185 (b) If the contribution is primarily a new model architecture, the paper should describe
186 the architecture clearly and fully.
 - 187 (c) If the contribution is a new model (e.g., a large language model), then there should
188 either be a way to access this model for reproducing the results or a way to reproduce
189 the model (e.g., with an open-source dataset or instructions for how to construct
190 the dataset).
 - 191 (d) We recognize that reproducibility may be tricky in some cases, in which case
192 authors are welcome to describe the particular way they provide for reproducibility.
193 In the case of closed-source models, it may be that access to the model is limited in
194 some way (e.g., to registered users), but it should be possible for other researchers
195 to have some path to reproducing or verifying the results.

196 5. Open access to data and code

197 Question: Does the paper provide open access to the data and code, with sufficient instruc-
198 tions to faithfully reproduce the main experimental results, as described in supplemental
199 material?

200 Answer: [Yes]

201 Justification: The data is openly available [6]. And the code is added as a footnote in section
202 2.3

203 Guidelines:

- 204 • The answer NA means that paper does not include experiments requiring code.
- 205 • Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 206 • While we encourage the release of code and data, we understand that this might not be
207 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not
208 including code, unless this is central to the contribution (e.g., for a new open-source
209 benchmark).
- 210 • The instructions should contain the exact command and environment needed to run to
211 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 212 • The authors should provide instructions on data access and preparation, including how
213 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 214 • The authors should provide scripts to reproduce all experimental results for the new
215 proposed method and baselines. If only a subset of experiments are reproducible, they
216 should state which ones are omitted from the script and why.
- 217 • At submission time, to preserve anonymity, the authors should release anonymized
218 versions (if applicable).
- 219 • Providing as much information as possible in supplemental material (appended to the
220 paper) is recommended, but including URLs to data and code is permitted.

223 6. Experimental setting/details

224 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-
225 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the
226 results?

227 Answer: [Yes]

228 Justification: The training and test splits are specified in the methodology.

229 Guidelines:

- 230 • The answer NA means that the paper does not include experiments.
- 231 • The experimental setting should be presented in the core of the paper to a level of detail
232 that is necessary to appreciate the results and make sense of them.
- 233 • The full details can be provided either with the code, in appendix, or as supplemental
234 material.

235 7. Experiment statistical significance

236 Question: Does the paper report error bars suitably and correctly defined or other appropriate
237 information about the statistical significance of the experiments?

238 Answer: [No]

239 Justification: Not enough experimental repetitions or folds were conducted to produce
240 statistically significant results or perform formal statistical tests. Reported performance
241 metrics therefore reflect central tendencies from the available runs rather than statistical
242 estimates.

243 Guidelines:

- 244 • The answer NA means that the paper does not include experiments.
- 245 • The authors should answer “Yes” if the results are accompanied by error bars, confi-
246 dence intervals, or statistical significance tests, at least for the experiments that support
247 the main claims of the paper.

248 • The factors of variability that the error bars are capturing should be clearly stated (for
 249 example, train/test split, initialization, random drawing of some parameter, or overall
 250 run with given experimental conditions).
 251 • The method for calculating the error bars should be explained (closed form formula,
 252 call to a library function, bootstrap, etc.)
 253 • The assumptions made should be given (e.g., Normally distributed errors).
 254 • It should be clear whether the error bar is the standard deviation or the standard error
 255 of the mean.
 256 • It is OK to report 1-sigma error bars, but one should state it. The authors should
 257 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
 258 of Normality of errors is not verified.
 259 • For asymmetric distributions, the authors should be careful not to show in tables or
 260 figures symmetric error bars that would yield results that are out of range (e.g. negative
 261 error rates).
 262 • If error bars are reported in tables or plots, The authors should explain in the text how
 263 they were calculated and reference the corresponding figures or tables in the text.

264 **8. Experiments compute resources**

265 Question: For each experiment, does the paper provide sufficient information on the com-
 266 puter resources (type of compute workers, memory, time of execution) needed to reproduce
 267 the experiments?

268 Answer: [Yes]

269 Justification: The compute used is specified in the methodology.

270 Guidelines:

271 • The answer NA means that the paper does not include experiments.
 272 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
 273 or cloud provider, including relevant memory and storage.
 274 • The paper should provide the amount of compute required for each of the individual
 275 experimental runs as well as estimate the total compute.
 276 • The paper should disclose whether the full research project required more compute
 277 than the experiments reported in the paper (e.g., preliminary or failed experiments that
 278 didn't make it into the paper).

279 **9. Code of ethics**

280 Question: Does the research conducted in the paper conform, in every respect, with the
 281 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

282 Answer: [Yes]

283 Justification: The code of ethics was reviewed, and this research conforms to it.

284 Guidelines:

285 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
 286 • If the authors answer No, they should explain the special circumstances that require a
 287 deviation from the Code of Ethics.
 288 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
 289 eration due to laws or regulations in their jurisdiction).

290 **10. Broader impacts**

291 Question: Does the paper discuss both potential positive societal impacts and negative
 292 societal impacts of the work performed?

293 Answer: [Yes]

294 Justification: The positive impact of reliable biomedical image segmentation is discussed in
 295 the introduction.

296 Guidelines:

297 • The answer NA means that there is no societal impact of the work performed.

- 298 • If the authors answer NA or No, they should explain why their work has no societal
299 impact or why the paper does not address societal impact.
- 300 • Examples of negative societal impacts include potential malicious or unintended uses
301 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
302 (e.g., deployment of technologies that could make decisions that unfairly impact specific
303 groups), privacy considerations, and security considerations.
- 304 • The conference expects that many papers will be foundational research and not tied
305 to particular applications, let alone deployments. However, if there is a direct path to
306 any negative applications, the authors should point it out. For example, it is legitimate
307 to point out that an improvement in the quality of generative models could be used to
308 generate deepfakes for disinformation. On the other hand, it is not needed to point out
309 that a generic algorithm for optimizing neural networks could enable people to train
310 models that generate Deepfakes faster.
- 311 • The authors should consider possible harms that could arise when the technology is
312 being used as intended and functioning correctly, harms that could arise when the
313 technology is being used as intended but gives incorrect results, and harms following
314 from (intentional or unintentional) misuse of the technology.
- 315 • If there are negative societal impacts, the authors could also discuss possible mitigation
316 strategies (e.g., gated release of models, providing defenses in addition to attacks,
317 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
318 feedback over time, improving the efficiency and accessibility of ML).

319 11. Safeguards

320 Question: Does the paper describe safeguards that have been put in place for responsible
321 release of data or models that have a high risk for misuse (e.g., pretrained language models,
322 image generators, or scraped datasets)?

323 Answer: [NA]

324 Justification: No new model or dataset with potential for misuse was released.

325 Guidelines:

- 326 • The answer NA means that the paper poses no such risks.
- 327 • Released models that have a high risk for misuse or dual-use should be released with
328 necessary safeguards to allow for controlled use of the model, for example by requiring
329 that users adhere to usage guidelines or restrictions to access the model or implementing
330 safety filters.
- 331 • Datasets that have been scraped from the Internet could pose safety risks. The authors
332 should describe how they avoided releasing unsafe images.
- 333 • We recognize that providing effective safeguards is challenging, and many papers do
334 not require this, but we encourage authors to take this into account and make a best
335 faith effort.

336 12. Licenses for existing assets

337 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
338 the paper, properly credited and are the license and terms of use explicitly mentioned and
339 properly respected?

340 Answer: [Yes]

341 Justification: Data and code sources are well cited.

342 Guidelines:

- 343 • The answer NA means that the paper does not use existing assets.
- 344 • The authors should cite the original paper that produced the code package or dataset.
- 345 • The authors should state which version of the asset is used and, if possible, include a
346 URL.
- 347 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 348 • For scraped data from a particular source (e.g., website), the copyright and terms of
349 service of that source should be provided.

350 • If assets are released, the license, copyright information, and terms of use in the
351 package should be provided. For popular datasets, paperswithcode.com/datasets
352 has curated licenses for some datasets. Their licensing guide can help determine the
353 license of a dataset.
354 • For existing datasets that are re-packaged, both the original license and the license of
355 the derived asset (if it has changed) should be provided.
356 • If this information is not available online, the authors are encouraged to reach out to
357 the asset's creators.

358 **13. New assets**

359 Question: Are new assets introduced in the paper well documented and is the documentation
360 provided alongside the assets?

361 Answer: **[Yes]**

362 Justification: The new assets introduced are the loss composite loss functions, available in
363 the provided code.

364 Guidelines:

365 • The answer NA means that the paper does not release new assets.
366 • Researchers should communicate the details of the dataset/code/model as part of their
367 submissions via structured templates. This includes details about training, license,
368 limitations, etc.
369 • The paper should discuss whether and how consent was obtained from people whose
370 asset is used.
371 • At submission time, remember to anonymize your assets (if applicable). You can either
372 create an anonymized URL or include an anonymized zip file.

373 **14. Crowdsourcing and research with human subjects**

374 Question: For crowdsourcing experiments and research with human subjects, does the paper
375 include the full text of instructions given to participants and screenshots, if applicable, as
376 well as details about compensation (if any)?

377 Answer: **[NA]**

378 Justification: No crowdsourcing was done for this work.

379 Guidelines:

380 • The answer NA means that the paper does not involve crowdsourcing nor research with
381 human subjects.
382 • Including this information in the supplemental material is fine, but if the main contribu-
383 tion of the paper involves human subjects, then as much detail as possible should be
384 included in the main paper.
385 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,
386 or other labor should be paid at least the minimum wage in the country of the data
387 collector.

388 **15. Institutional review board (IRB) approvals or equivalent for research with human
389 subjects**

390 Question: Does the paper describe potential risks incurred by study participants, whether
391 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
392 approvals (or an equivalent approval/review based on the requirements of your country or
393 institution) were obtained?

394 Answer: **[NA]**

395 Justification: No human subjects research was conducted

396 Guidelines:

397 • The answer NA means that the paper does not involve crowdsourcing nor research with
398 human subjects.
399 • Depending on the country in which research is conducted, IRB approval (or equivalent)
400 may be required for any human subjects research. If you obtained IRB approval, you
401 should clearly state this in the paper.

402 • We recognize that the procedures for this may vary significantly between institutions
 403 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
 404 guidelines for their institution.
 405 • For initial submissions, do not include any information that would break anonymity (if
 406 applicable), such as the institution conducting the review.

407 **16. Declaration of LLM usage**

408 Question: Does the paper describe the usage of LLMs if it is an important, original, or
 409 non-standard component of the core methods in this research? Note that if the LLM is used
 410 only for writing, editing, or formatting purposes and does not impact the core methodology,
 411 scientific rigorosity, or originality of the research, declaration is not required.

412 Answer: **[No]**

413 Justification: No large language model was used in the design, analysis, or generation of
 414 results; only standard writing assistance was used.

415 Guidelines:

416 • The answer NA means that the core method development in this research does not
 417 involve LLMs as any important, original, or non-standard components.
 418 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)
 419 for what should or should not be described.

420 **References**

421 [1] Khadijeh Askaripour and Arkadiusz Zak. Breast mri segmentation by deep learning: Key gaps
 422 and challenges. *IEEE Access*, 11:117935–117946, 2023.

423 [2] Nathaniel M Braman, Maryam Etesami, Prateek Prasanna, Christina Dubchuk, Hannah Gilmore,
 424 Pallavi Tiwari, Donna Plecha, and Anant Madabhushi. Intratumoral and peritumoral radiomics
 425 for the pretreatment prediction of pathological complete response to neoadjuvant chemotherapy
 426 based on breast dce-mri. *Breast Cancer Research*, 19:1–14, 2017.

427 [3] Honghao Chen, Xiangxiang Chu, Yongjian Ren, Xin Zhao, and Kaiqi Huang. Pelk: Parameter-
 428 efficient large kernel convnets with peripheral convolution. In *Proceedings of the IEEE/CVF
 429 conference on computer vision and pattern recognition*, pages 5557–5567, 2024.

430 [4] Kawtar Debbi, Paul Habert, Anaïs Grob, Anderson Loundou, Pascale Siles, Axel Bartoli, and
 431 Alexis Jacquier. Radiomics model to classify mammary masses using breast dce-mri compared
 432 to the bi-rads classification performance. *Insights into Imaging*, 14(1):64, 2023.

433 [5] Khuram Faraz, Grégoire Dauce, Amine Bouhamama, Benjamin Leporq, Hajime Sasaki, Yoshi-
 434 taka Bito, Olivier Beuf, and Frank Pilleul. Characterization of breast tumors from mr images us-
 435 ing radiomics and machine learning approaches. *Journal of Personalized Medicine*, 13(7):1062,
 436 2023.

437 [6] Lidia Garrucho, Kaisar Kushibar, Claire-Anne Reidel, Smriti Joshi, Richard Osuala, Apos-
 438 tolia Tsirikoglou, Maciej Bobowicz, Javier Del Riego, Alessandro Catanese, Katarzyna
 439 Gwoźdiewicz, et al. A large-scale multicenter breast cancer dce-mri benchmark dataset
 440 with expert segmentations. *Scientific data*, 12(1):453, 2025.

441 [7] Briete Goorts, Kelly MA Dreuning, Janneke B Houwers, Loes FS Kooreman, Evert-Jan G
 442 Boerma, Ritse M Mann, Marc BI Lobbes, and Marjolein L Smidt. Mri-based response patterns
 443 during neoadjuvant chemotherapy can predict pathological (complete) response in patients with
 444 breast cancer. *Breast Cancer Research*, 20(1):34, 2018.

445 [8] Gabriel N Hortobagyi, Jaime de la Garza Salazar, Kathleen Pritchard, Dino Amadori, Renate
 446 Haidinger, Clifford A Hudis, Hussein Khaled, Mei-Ching Liu, Miguel Martin, Moise Namer,
 447 et al. The global breast cancer burden: variations in epidemiology and survival. *Clinical breast
 448 cancer*, 6(5):391–401, 2005.

449 [9] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net:
 450 a self-configuring method for deep learning-based biomedical image segmentation. *Nature
 451 methods*, 18(2):203–211, 2021.

452 [10] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining
453 Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision*
454 and pattern recognition, pages 11976–11986, 2022.

455 [11] Toufiq Musah. Large kernel mednext for breast tumor segmentation and self-normalizing
456 network for pcr classification in magnetic resonance images. In *Deep Breast Workshop on AI*
457 and Imaging for Diagnostic and Treatment Challenges in Breast Care, pages 72–80. Springer,
458 2025.

459 [12] Toufiq Musah, Chinasa Kalaiwo, Maimoona Akram, Ubaida Napari Abdulai, Maruf Adewole,
460 Farouk Dako, Adaobi Chiaazor Emegoakor, Uduonna C Anazodo, Prince Ebenezer Adjei, and
461 Confidence Raymond. Towards trustworthy breast tumor segmentation in ultrasound using
462 monte carlo dropout and deep ensembles for epistemic uncertainty estimation. *arXiv preprint*
463 *arXiv:2508.17768*, 2025.

464 [13] JPM O'Donnell, SA Gasior, MG Davey, E O'Malley, AJ Lowery, J McGarry, AM O'Connell,
465 MJ Kerin, and P McCarthy. The accuracy of breast mri radiomic methodologies in predicting
466 pathological complete response to neoadjuvant chemotherapy: A systematic review and network
467 meta-analysis. *European journal of radiology*, 157:110561, 2022.

468 [14] Selvi Radhakrishna, S Agarwal, Purvish M Parikh, K Kaur, Shikha Panwar, Shelly Sharma,
469 Ashish Dey, KK Saxena, Madhavi Chandra, and Seema Sud. Role of magnetic resonance
470 imaging in breast cancer management. *South Asian journal of cancer*, 7(2):69, 2018.

471 [15] Saikat Roy, Gregor Koehler, Constantin Ulrich, Michael Baumgartner, Jens Petersen, Fabian
472 Isensee, Paul F Jaeger, and Klaus H Maier-Hein. Mednext: transformer-driven scaling of
473 convnets for medical image segmentation. In *International Conference on Medical Image*
474 *Computing and Computer-Assisted Intervention*, pages 405–415. Springer, 2023.

475 [16] Guoliang Shao, Linyin Fan, Juan Zhang, Gang Dai, and Tieming Xie. Association of dw/dce-
476 mri features with prognostic factors in breast cancer. *The International journal of biological*
477 *markers*, 32(1):118–125, 2017.

478 [17] Joost JM Van Griethuysen, Andriy Fedorov, Chintan Parmar, Ahmed Hosny, Nicole Au-
479 coin, Vivek Narayan, Regina GH Beets-Tan, Jean-Christophe Fillion-Robin, Steve Pieper, and
480 Hugo JW Aerts. Computational radiomics system to decode the radiographic phenotype.
481 *Cancer research*, 77(21):e104–e107, 2017.