
AI-Driven Demand-Oriented STEM Education Strategy for Our Muslim Community

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In the context of rapidly advancing global technology, Science, Technology, En-
2 gineering, and Mathematics (STEM) education in the Middle East and Muslim-
3 majority regions is essential for driving innovation and supporting economic diver-
4 sification. But, significant gaps remain between current educational practices and
5 our vision, particularly in learning methodologies, student motivation, employment
6 market alignment, and educational equity. We propose a comprehensive strategic
7 framework that leverages large language models (LLMs) and virtual reality (VR) to
8 create an AI-supported, closed-loop skills training system offering immersive and
9 personalized learning experiences. Additionally, it promotes a mutually beneficial,
10 cross-regional educational cooperation model that fosters resource sharing between
11 economically developed and underdeveloped areas to support the development
12 of Muslim communities around the world. This framework aims to establish an
13 inclusive and efficient global STEM education system within Muslim communi-
14 ties, empowering the younger generation to meet future challenges while ensuring
15 sustainable returns for sponsors in cultivating global tech talent.

16 1 Introduction

17 STEM education in the global Muslim community is tasked with enhancing innovation capacity,
18 cultivating globally competitive talent, and supporting economic diversification. But, gaps between
19 educational content and practical needs, limited adaptability of curricula to cultural contexts, and
20 imbalanced resource allocation hinder the achievement of these goals. We propose a multifaceted
21 strategic approach to bridge the gap between existing educational systems and the region's economic
22 and social vision, as shown in figure 1. Key strategies include the introduction of LLMs trained in
23 specialized fields to enable dynamic integration of cross-disciplinary knowledge, creating an education
24 system that meets contemporary demands; utilizing VR-driven experiential learning to enhance
25 students' practical skills and engagement; establishing an AI-supported, closed-loop training system
26 to align skill development precisely with job market needs; and fostering a mutually beneficial, cross-
27 regional educational cooperation model to facilitate resource sharing and talent cultivation within
28 Muslim communities, promoting collaboration between economically developed and underdeveloped
29 areas. This framework aims to integrate innovative technology and educational resources to help
30 Muslim regions address the dual challenges of educational and economic transformation [1].

31 2 LLM-Driven Education: Interdisciplinary Collaboration and Resource 32 Equity

33 With the advancement of AI, the potential of LLMs in education is gradually emerging. Through
34 the dynamic integration of knowledge across various fields, intelligent adaptive learning features,

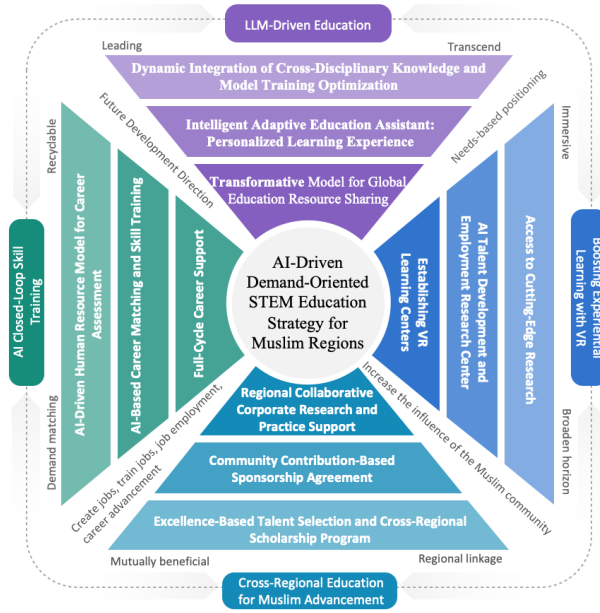


Figure 1: AI-Driven Demand-Oriented STEM Education Strategic Framework for Muslim Regions

35 and transformative solutions for global educational resource sharing, LLMs offer new pathways to
 36 address educational challenges [2].

37 **2.1 Dynamic Integration of Cross-Disciplinary Knowledge and Model Training Optimization**

38 The effectiveness of LLMs lies in their breadth of knowledge and timely updates. To ensure
 39 high reliability, experts from various fields continuously update the model’s knowledge base to
 40 incorporate the latest advancements in areas such as physics, artificial intelligence, and medicine.
 41 This comprehensive training process not only enhances the precision of knowledge output but also
 42 promotes a more balanced distribution of educational resources, breaking the regional limitations of
 43 subject-specific resources.

44 **2.2 Intelligent Adaptive Education Assistant: Personalized Learning Experience**

45 LLMs possess adaptive capabilities that allow them to adjust content and difficulty levels in real-
 46 time based on learner feedback, providing personalized learning support. The model can generate
 47 multi-level knowledge modules to meet the needs of students at different learning levels. This
 48 intelligent adaptability not only facilitates individualized learning but also compensates for traditional
 49 education’s limitations in personalization, thereby enhancing students’ learning engagement [3].

50 **2.3 Global Education Resource Sharing: A New Direction**

51 Although the application of LLMs in education is still in its experimental stages, their potential to
 52 address global educational inequalities and delays in knowledge updates is increasingly evident. As a
 53 forward-looking educational tool, LLMs offer innovative solutions for global education. Developed
 54 nations are piloting these models in classroom teaching and personalized learning to enhance interac-
 55 tivity and efficiency, as shown in figure 2, while in regions with limited educational resources, the
 56 model provides an alternative learning support. Studies suggest that this approach will significantly
 57 improve global educational accessibility, heralding profound changes in the global education system.

Comparative Evaluation of AI-Driven and Traditional Learning Systems Across Key Educational Dimensions

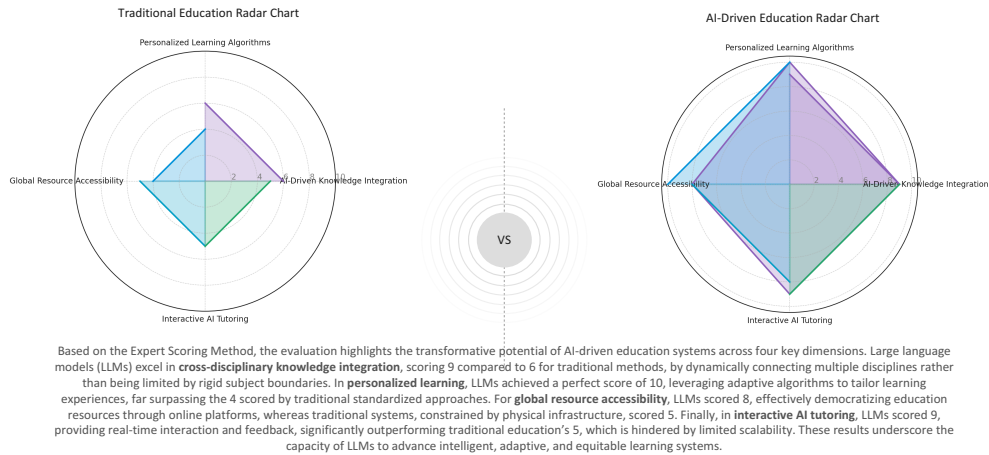


Figure 2: Comparison of AI-Driven and Traditional Learning Systems in Key Educational Dimensions

58 3 Enhancing Experiential Learning through Virtual Reality

59 VR technology has brought innovation to education by offering students immersive learning experiences, allowing them to gain a deep understanding of practical applications within STEM fields.
 60 The application of VR not only sparks students' interest in learning but also enables them to explore complex scientific concepts within a virtual environment, laying a solid foundation for their future career development.
 61
 62
 63

64 3.1 Access to Cutting-Edge Research

65 Through VR simulations, students can virtually visit international research laboratories, explore space, or participate in biomedical experiments. These experiences help students understand the societal impact of science and technology, increase their motivation to learn, and inspire them to explore STEM-related careers [4-5].
 66
 67
 68

69 3.2 Establishing VR Learning Centers

70 Educational institutions in the Muslim countries can collaborate with government and industry to establish VR learning centers. These centers not only provide academic knowledge but also integrate practical experience into immersive courses, offering a replicable model for advancing regional STEM education.
 71
 72
 73

74 3.3 AI Talent Development and Employment Research Center

75 Through a dedicated AI research center, career pathways can be analyzed and aligned with industry needs. This center sets specific development goals and designs targeted training programs to prepare students for real job opportunities, creating a seamless "learning-to-employment" pathway. The curriculum is closely aligned with labor market demands and integrates feedback from both government and industry.
 76
 77
 78
 79

80 **4 AI-Based Closed-Loop Training System for Skill Matching**

81 The disconnect between education and labor market demands is a key challenge in talent development.
82 The AI-supported closed-loop training system uses intelligent assessment and skill matching to ensure
83 that students acquire practical skills that are directly applicable to the job market.

84 **4.1 AI-Driven Human Resource Model for Career Assessment**

85 The AI system conducts a comprehensive assessment of students based on behavioral traits and
86 cognitive abilities to understand their skills and career inclinations. Through big data analysis, AI
87 can create personalized training plans that align students' potential with specific career roles, thereby
88 enhancing the effectiveness of education and increasing employment success rates.

89 **4.2 AI-Based Career Matching and Skill Training**

90 After assessing students' career directions, the AI system matches them with specific training
91 programs to help them acquire skills directly required by the job market. Additionally, the AI
92 dynamically adjusts training content based on feedback to ensure that the curriculum aligns with
93 industry needs. Graduates can also receive job matching, mentorship, and career development advice
94 through the AI system, facilitating a smooth transition from learning to employment.

95 **4.3 Full-Cycle Career Support**

96 The AI system provides students with full-cycle support from initial training to career development.
97 By continuously monitoring learning and job performance, the AI can recommend additional skill
98 training and advancement opportunities, helping students adapt to the work environment and gradually
99 build career confidence, laying the foundation for further professional growth.

100 **5 Mutually Beneficial Educational Collaboration Model for Muslim 101 Communities**

102 Promoting educational resource sharing and cultivating local high-skilled talent are key to enhanc-
103 ing the overall competitiveness of the Muslim community. Through cross-regional collaboration,
104 research institutions and high-tech enterprises from economically developed countries can support
105 underdeveloped regions, while building global technology and talent reserves to achieve mutual
106 benefits in education and economic growth.

107 **5.1 Excellence-Based Talent Selection and Cross-Regional Scholarship Program**

108 Establish scholarships funded by research institutions or enterprises from developed countries to select
109 outstanding students from underdeveloped regions within Muslim communities for admission to top
110 universities. This not only cultivates future tech talent but also helps sponsors attract high-potential
111 international talent and strengthen global R&D teams.

112 **5.2 Community Contribution-Based Sponsorship Agreement**

113 Through a contribution-based sponsorship agreement, sponsored students are required to serve the
114 Muslim community for a certain number of years upon completing their studies. This ensures that
115 the educational investment translates into long-term development for the community and provides
116 sponsors with a foundation for expanding their business and influence in local markets, fostering a
117 sustainable and mutually beneficial partnership.

118 **5.3 Regional Collaborative Corporate Research and Practice Support**

119 Encourage high-tech companies to collaborate with universities and research institutions in the
120 Muslim community to establish cross-border research internships and training programs. Participating
121 companies can not only cultivate technical talent that meets their needs but also expand their business
122 and brand influence in Muslim markets, laying a solid foundation for global outreach.

123 **References**

- 124 [1] Zhao, Q.P. (2020) Integration and application of virtual reality and artificial intelligence in education. In
125 *VR+AI: The Potential Ultimate Educational Technology*. Beijing, China: Institute of Smart Learning, Beijing
126 Normal University.
- 127 [2] School of Information Science and Technology, Beijing Normal University (2020) Artificial intelligence in
128 personalized STEM education. In *Research on AI Application in Personalized STEM Education*. Beijing, China:
129 China Educational Informatization.
- 130 [3] Wang, Y.F. (2024) The strategic value of large scientific facilities and their position in the science and
131 technology innovation system. In *Science and Technology Daily*. Beijing, China: Science and Technology Daily
132 Press.
- 133 [4] Khan, A.S. (2023) The strategic role of big science in national development. In *Journal of Science Policy*.
134 Cambridge, MA: Elsevier.
- 135 [5] UNESCO (2022) Research and development in the Arab states. In *UNESCO Report*. Paris, France: UNESCO
136 Publishing.

137 **NeurIPS Paper Checklist**

138 **1. Claims**

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

141 Answer: [\[Yes\]](#) ,

142 Justification: The claims in the abstract and introduction accurately reflect the contributions
143 of the paper, specifically highlighting the novel application of deep learning models in the
144 analysis of biophysics data.

145 Guidelines:

- 146 • The answer NA means that the abstract and introduction do not include the claims
147 made in the paper.
- 148 • The abstract and/or introduction should clearly state the claims made, including the
149 contributions made in the paper and important assumptions and limitations. A No or
150 NA answer to this question will not be perceived well by the reviewers.
- 151 • The claims made should match theoretical and experimental results, and reflect how
152 much the results can be expected to generalize to other settings.
- 153 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
154 are not attained by the paper.

155 **2. Limitations**

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

157 Answer: [\[Yes\]](#)

158 Justification: The paper includes a dedicated section on limitations (Section 4). We address
159 the limitations related to the generalizability of our approach, as the results have been
160 validated on a limited number of datasets.

161 Guidelines:

- 162 • The answer NA means that the paper has no limitation while the answer No means that
163 the paper has limitations, but those are not discussed in the paper.
- 164 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 165 • The paper should point out any strong assumptions and how robust the results are to
166 violations of these assumptions (e.g., independence assumptions, noiseless settings,
167 model well-specification, asymptotic approximations only holding locally). The authors
168 should reflect on how these assumptions might be violated in practice and what the
169 implications would be.
- 170 • The authors should reflect on the scope of the claims made, e.g., if the approach was
171 only tested on a few datasets or with a few runs. In general, empirical results often
172 depend on implicit assumptions, which should be articulated.
- 173 • The authors should reflect on the factors that influence the performance of the approach.
174 For example, a facial recognition algorithm may perform poorly when image resolution
175 is low or images are taken in low lighting. Or a speech-to-text system might not be
176 used reliably to provide closed captions for online lectures because it fails to handle
177 technical jargon.
- 178 • The authors should discuss the computational efficiency of the proposed algorithms
179 and how they scale with dataset size.
- 180 • If applicable, the authors should discuss possible limitations of their approach to
181 address problems of privacy and fairness.
- 182 • While the authors might fear that complete honesty about limitations might be used by
183 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
184 limitations that aren't acknowledged in the paper. The authors should use their best
185 judgment and recognize that individual actions in favor of transparency play an impor-
186 tant role in developing norms that preserve the integrity of the community. Reviewers
187 will be specifically instructed to not penalize honesty concerning limitations.

188 **3. Theory Assumptions and Proofs**

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

191 Answer: [NA]

192 Justification: This paper focuses on the application of deep learning models for 3D re-
193 construction and does not introduce any new theoretical results that require formal proofs.
194 The methods and algorithms used are based on existing techniques, which are referenced
195 appropriately.

196 Guidelines:

- 197 • The answer NA means that the paper does not include theoretical results.
- 198 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
199 referenced.
- 200 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 201 • The proofs can either appear in the main paper or the supplemental material, but if
202 they appear in the supplemental material, the authors are encouraged to provide a short
203 proof sketch to provide intuition.
- 204 • Inversely, any informal proof provided in the core of the paper should be complemented
205 by formal proofs provided in appendix or supplemental material.
- 206 • Theorems and Lemmas that the proof relies upon should be properly referenced.

207 4. Experimental Result Reproducibility

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

211 Answer: [Yes]

212 Justification: The paper provides detailed descriptions of the experimental setup, data
213 acquisition methods, and model architectures, including the use of super-resolution and
214 diffusion models. Additionally, all code, model training scripts, and datasets have been made
215 available on GitHub (as mentioned in Section 3.3), allowing for full reproducibility of the
216 results. Furthermore, the paper outlines the specific hardware and software configurations
217 used for training.

218 Guidelines:

- 219 • The answer NA means that the paper does not include experiments.
- 220 • If the paper includes experiments, a No answer to this question will not be perceived
221 well by the reviewers: Making the paper reproducible is important, regardless of
222 whether the code and data are provided or not.
- 223 • If the contribution is a dataset and/or model, the authors should describe the steps taken
224 to make their results reproducible or verifiable.
- 225 • Depending on the contribution, reproducibility can be accomplished in various ways.
226 For example, if the contribution is a novel architecture, describing the architecture fully
227 might suffice, or if the contribution is a specific model and empirical evaluation, it may
228 be necessary to either make it possible for others to replicate the model with the same
229 dataset, or provide access to the model. In general, releasing code and data is often
230 one good way to accomplish this, but reproducibility can also be provided via detailed
231 instructions for how to replicate the results, access to a hosted model (e.g., in the case
232 of a large language model), releasing of a model checkpoint, or other means that are
233 appropriate to the research performed.
- 234 • While NeurIPS does not require releasing code, the conference does require all submis-
235 sions to provide some reasonable avenue for reproducibility, which may depend on the
236 nature of the contribution. For example
 - 237 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
238 to reproduce that algorithm.
 - 239 (b) If the contribution is primarily a new model architecture, the paper should describe
240 the architecture clearly and fully.

- 241 (c) If the contribution is a new model (e.g., a large language model), then there should
242 either be a way to access this model for reproducing the results or a way to reproduce
243 the model (e.g., with an open-source dataset or instructions for how to construct
244 the dataset).
- 245 (d) We recognize that reproducibility may be tricky in some cases, in which case
246 authors are welcome to describe the particular way they provide for reproducibility.
247 In the case of closed-source models, it may be that access to the model is limited in
248 some way (e.g., to registered users), but it should be possible for other researchers
249 to have some path to reproducing or verifying the results.

250 5. Open access to data and code

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

254 Answer: [Yes]

255 Justification: The paper provides open access to both the code and datasets through a publicly
256 accessible GitHub repository, as stated in Section 3.3. Detailed instructions are included for
257 data access, environment setup, and running the code, ensuring that the experiments can be
258 faithfully reproduced.

259 Guidelines:

- 260 • The answer NA means that paper does not include experiments requiring code.
- 261 • Please see the NeurIPS code and data submission guidelines ([https://nips.cc/
262 public/guides/CodeSubmissionPolicy](https://nips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 263 • While we encourage the release of code and data, we understand that this might not be
264 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not
265 including code, unless this is central to the contribution (e.g., for a new open-source
266 benchmark).
- 267 • The instructions should contain the exact command and environment needed to run to
268 reproduce the results. See the NeurIPS code and data submission guidelines ([https://
269 nips.cc/public/guides/CodeSubmissionPolicy](https://nips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 270 • The authors should provide instructions on data access and preparation, including how
271 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 272 • The authors should provide scripts to reproduce all experimental results for the new
273 proposed method and baselines. If only a subset of experiments are reproducible, they
274 should state which ones are omitted from the script and why.
- 275 • At submission time, to preserve anonymity, the authors should release anonymized
276 versions (if applicable).
- 277 • Providing as much information as possible in supplemental material (appended to the
278 paper) is recommended, but including URLs to data and code is permitted.

279 6. Experimental Setting/Details

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

283 Answer: [Yes]

284 Justification: The paper specifies the details of the experimental setup, including data acquisi-
285 tion, data splits, and model hyperparameters. The choice of optimizer and its configurations
286 are also provided, ensuring that the experimental results can be fully understood. Further
287 details are provided in the supplemental material and the accompanying GitHub repository
288 (Section 2.1 and 2.2).”

289 Guidelines:

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

295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346

7. Experiment Statistical Significance

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

Answer: [No]

Justification: The paper focuses on the qualitative performance of the Vision-Cell framework in sparse data environments and does not include statistical significance measures such as error bars. The experiments are designed to demonstrate the framework's ability to generate high-quality 3D reconstructions rather than to assess variability or uncertainty across multiple runs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

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

Answer: [Yes]

Justification: The paper specifies the type of compute resources used, including an RTX 4060 GPU and later two A100 GPUs for more complex model training. Memory limitations encountered with the RTX 4060 GPU are also discussed, along with the time required for training tasks. Further details, such as memory usage and time of execution, are provided in the supplemental material (see Section 4).

Guidelines:

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

9. Code Of Ethics

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

347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399

Answer: [Yes]

Justification: The research conducted in this paper fully conforms with the NeurIPS Code of Ethics. All experiments and data used in this study are from publicly available sources or obtained through ethical means, with no human or animal subjects involved. The paper adheres to the principles of transparency, fairness, and societal responsibility outlined in the Code of Ethics.

Guidelines:

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

10. Broader Impacts

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

Answer: [Yes]

Justification: The paper discusses the positive societal impacts of the Vision-Cell framework, particularly in its potential to advance biological and medical research by enabling high-precision 3D cell reconstructions with minimal physical damage to samples. This could lead to breakthroughs in fields such as neuroscience and structural biology. The paper also considers the broader implications of deep learning in biological research but does not foresee any direct negative societal impacts, as the work is foundational and not tied to applications with ethical risks.

Guidelines:

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

11. Safeguards

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

Answer: [NA]

Justification: The paper does not release any data or models that pose a high risk for misuse. The Vision-Cell framework focuses on biological data with no foreseeable ethical concerns.

400 related to privacy, security, or malicious use. Therefore, no specific safeguards are necessary
401 for the release of this work.

402 Guidelines:

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

413 12. Licenses for existing assets

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

417 Answer: [\[Yes\]](#)

418 Justification: The paper uses publicly available datasets and models, all of which are properly
419 credited in the references section. We have explicitly mentioned the license and terms of
420 use for these assets where applicable (see Section X for details), and no scraped data from
421 restricted sources were used

422 Guidelines:

- 423 • The answer NA means that the paper does not use existing assets.
- 424 • The authors should cite the original paper that produced the code package or dataset.
- 425 • The authors should state which version of the asset is used and, if possible, include a
426 URL.
- 427 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 428 • For scraped data from a particular source (e.g., website), the copyright and terms of
429 service of that source should be provided.
- 430 • If assets are released, the license, copyright information, and terms of use in the
431 package should be provided. For popular datasets, `paperswithcode.com/datasets`
432 has curated licenses for some datasets. Their licensing guide can help determine the
433 license of a dataset.
- 434 • For existing datasets that are re-packaged, both the original license and the license of
435 the derived asset (if it has changed) should be provided.
- 436 • If this information is not available online, the authors are encouraged to reach out to
437 the asset's creators.

438 13. New Assets

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

441 Answer: [\[Yes\]](#)

442 Justification: The Vision-Cell framework introduced in this paper is a new asset, and detailed
443 documentation is provided alongside the code and models, including instructions for usage,
444 training, and evaluation. The assets are released via a GitHub repository with a structured
445 template to ensure easy reproducibility (see Section 3.3).

446 Guidelines:

- 447 • The answer NA means that the paper does not release new assets.
- 448 • Researchers should communicate the details of the dataset/code/model as part of their
449 submissions via structured templates. This includes details about training, license,
450 limitations, etc.
- 451 • The paper should discuss whether and how consent was obtained from people whose
452 asset is used.

- 453 • At submission time, remember to anonymize your assets (if applicable). You can either
454 create an anonymized URL or include an anonymized zip file.

455 **14. Crowdsourcing and Research with Human Subjects**

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

459 Answer: [No] ,

460 Justification: The paper does not involve crowdsourcing or research with human subjects, as
461 it focuses on computational experiments using biological data obtained through non-human
462 means

463 Guidelines:

- 464 • The answer NA means that the paper does not involve crowdsourcing nor research with
465 human subjects.
- 466 • Including this information in the supplemental material is fine, but if the main contribu-
467 tion of the paper involves human subjects, then as much detail as possible should be
468 included in the main paper.
- 469 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,
470 or other labor should be paid at least the minimum wage in the country of the data
471 collector.

472 **15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human**
473 **Subjects**

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

478 Answer: [No]

479 Justification: The paper does not involve research with human subjects, and therefore, IRB
480 approval or equivalent is not applicable.

481 Guidelines:

- 482 • The answer NA means that the paper does not involve crowdsourcing nor research with
483 human subjects.
- 484 • Depending on the country in which research is conducted, IRB approval (or equivalent)
485 may be required for any human subjects research. If you obtained IRB approval, you
486 should clearly state this in the paper.
- 487 • We recognize that the procedures for this may vary significantly between institutions
488 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
489 guidelines for their institution.
- 490 • For initial submissions, do not include any information that would break anonymity (if
491 applicable), such as the institution conducting the review.