
Formatting Instructions For NeurIPS 2024

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The abstract paragraph should be indented $\frac{1}{2}$ inch (3 picas) on both the left- and
2 right-hand margins. Use 10 point type, with a vertical spacing (leading) of 11 points.
3 The word **Abstract** must be centered, bold, and in point size 12. Two line spaces
4 precede the abstract. The abstract must be limited to one paragraph.

5 1 Submission of papers to NeurIPS 2024

6 Please read the instructions below carefully and follow them faithfully.

7 1.1 Style

8 Papers to be submitted to NeurIPS 2024 must be prepared according to the instructions presented
9 here. Papers may only be up to **nine** pages long, including figures. Additional pages *containing only*
10 *acknowledgments and references* are allowed. Papers that exceed the page limit will not be reviewed,
11 or in any other way considered for presentation at the conference.

12 The margins in 2024 are the same as those in previous years.

13 Authors are required to use the NeurIPS L^AT_EX style files obtainable at the NeurIPS website as
14 indicated below. Please make sure you use the current files and not previous versions. Tweaking the
15 style files may be grounds for rejection.

16 1.2 Retrieval of style files

17 The style files for NeurIPS and other conference information are available on the website at

18 <http://www.neurips.cc/>

19 The file `neurips_2024.pdf` contains these instructions and illustrates the various formatting re-
20 quirements your NeurIPS paper must satisfy.

21 The only supported style file for NeurIPS 2024 is `neurips_2024.sty`, rewritten for L^AT_EX 2_ε.
22 **Previous style files for L^AT_EX 2.09, Microsoft Word, and RTF are no longer supported!**

23 The L^AT_EX style file contains three optional arguments: `final`, which creates a camera-ready copy,
24 `preprint`, which creates a preprint for submission to, e.g., arXiv, and `nonatbib`, which will not
25 load the `natbib` package for you in case of package clash.

26 **Preprint option** If you wish to post a preprint of your work online, e.g., on arXiv, using the
27 NeurIPS style, please use the `preprint` option. This will create a nonanonymized version of your
28 work with the text “Preprint. Work in progress.” in the footer. This version may be distributed as you
29 see fit, as long as you do not say which conference it was submitted to. Please **do not** use the `final`
30 option, which should **only** be used for papers accepted to NeurIPS.

31 At submission time, please omit the `final` and `preprint` options. This will anonymize your
32 submission and add line numbers to aid review. Please do *not* refer to these line numbers in your
33 paper as they will be removed during generation of camera-ready copies.

34 The file `neurips_2024.tex` may be used as a “shell” for writing your paper. All you have to do is
35 replace the author, title, abstract, and text of the paper with your own.

36 The formatting instructions contained in these style files are summarized in Sections 2, 3, and 4
37 below.

38 **2 General formatting instructions**

39 The text must be confined within a rectangle 5.5 inches (33 picas) wide and 9 inches (54 picas) long.
40 The left margin is 1.5 inch (9 picas). Use 10 point type with a vertical spacing (leading) of 11 points.
41 Times New Roman is the preferred typeface throughout, and will be selected for you by default.
42 Paragraphs are separated by $\frac{1}{2}$ line space (5.5 points), with no indentation.

43 The paper title should be 17 point, initial caps/lower case, bold, centered between two horizontal
44 rules. The top rule should be 4 points thick and the bottom rule should be 1 point thick. Allow $\frac{1}{4}$ inch
45 space above and below the title to rules. All pages should start at 1 inch (6 picas) from the top of the
46 page.

47 For the final version, authors’ names are set in boldface, and each name is centered above the
48 corresponding address. The lead author’s name is to be listed first (left-most), and the co-authors’
49 names (if different address) are set to follow. If there is only one co-author, list both author and
50 co-author side by side.

51 Please pay special attention to the instructions in Section 4 regarding figures, tables, acknowledgments,
52 and references.

53 **3 Headings: first level**

54 All headings should be lower case (except for first word and proper nouns), flush left, and bold.

55 First-level headings should be in 12-point type.

56 **3.1 Headings: second level**

57 Second-level headings should be in 10-point type.

58 **3.1.1 Headings: third level**

59 Third-level headings should be in 10-point type.

60 **Paragraphs** There is also a `\paragraph` command available, which sets the heading in bold, flush
61 left, and inline with the text, with the heading followed by 1 em of space.

62 **4 Citations, figures, tables, references**

63 These instructions apply to everyone.

64 **4.1 Citations within the text**

65 The `natbib` package will be loaded for you by default. Citations may be author/year or numeric, as
66 long as you maintain internal consistency. As to the format of the references themselves, any style is
67 acceptable as long as it is used consistently.

68 The documentation for `natbib` may be found at

69 <http://mirrors.ctan.org/macros/latex/contrib/natbib/natnotes.pdf>




Figure 1: Sample figure caption.

70 Of note is the command `\citet`, which produces citations appropriate for use in inline text. For
71 example,

```
72 \citet{hasselmo} investigated\dots
```

73 produces

```
74 Hasselmo, et al. (1995) investigated...
```

75 If you wish to load the `natbib` package with options, you may add the following before loading the
76 `neurips_2024` package:

```
77 \PassOptionsToPackage{options}{natbib}
```

78 If `natbib` clashes with another package you load, you can add the optional argument `nonatbib`
79 when loading the style file:

```
80 \usepackage[nonatbib]{neurips_2024}
```

81 As submission is double blind, refer to your own published work in the third person. That is, use “In
82 the previous work of Jones et al. [4],” not “In our previous work [4].” If you cite your other papers
83 that are not widely available (e.g., a journal paper under review), use anonymous author names in the
84 citation, e.g., an author of the form “A. Anonymous” and include a copy of the anonymized paper in
85 the supplementary material.

86 4.2 Footnotes

87 Footnotes should be used sparingly. If you do require a footnote, indicate footnotes with a number¹
88 in the text. Place the footnotes at the bottom of the page on which they appear. Precede the footnote
89 with a horizontal rule of 2 inches (12 picas).

90 Note that footnotes are properly typeset *after* punctuation marks.²

91 4.3 Figures

92 All artwork must be neat, clean, and legible. Lines should be dark enough for purposes of reproduction.
93 The figure number and caption always appear after the figure. Place one line space before the figure
94 caption and one line space after the figure. The figure caption should be lower case (except for first
95 word and proper nouns); figures are numbered consecutively.

96 You may use color figures. However, it is best for the figure captions and the paper body to be legible
97 if the paper is printed in either black/white or in color.

¹Sample of the first footnote.

²As in this example.

Table 1: Sample table title

Part		
Name	Description	Size (μm)
Dendrite	Input terminal	~ 100
Axon	Output terminal	~ 10
Soma	Cell body	up to 10^6

98 **4.4 Tables**

99 All tables must be centered, neat, clean and legible. The table number and title always appear before
100 the table. See Table 1.

101 Place one line space before the table title, one line space after the table title, and one line space after
102 the table. The table title must be lower case (except for first word and proper nouns); tables are
103 numbered consecutively.

104 Note that publication-quality tables *do not contain vertical rules*. We strongly suggest the use of the
105 `booktabs` package, which allows for typesetting high-quality, professional tables:

106 `https://www.ctan.org/pkg/booktabs`

107 This package was used to typeset Table 1.

108 **4.5 Math**

109 Note that display math in bare TeX commands will not create correct line numbers for sub-
110 mission. Please use LaTeX (or AMSTeX) commands for unnumbered display math. (You
111 really shouldn't be using $\$$ anyway; see <https://tex.stackexchange.com/questions/503/why-is-preferable-to> and
112 <https://tex.stackexchange.com/questions/40492/what-are-the-differences-between-align-equation-and-displaymath> for more infor-
113 mation.)
114

115 **4.6 Final instructions**

116 Do not change any aspects of the formatting parameters in the style files. In particular, do not modify
117 the width or length of the rectangle the text should fit into, and do not change font sizes (except
118 perhaps in the **References** section; see below). Please note that pages should be numbered.

119 **5 Preparing PDF files**

120 Please prepare submission files with paper size "US Letter," and not, for example, "A4."

121 Fonts were the main cause of problems in the past years. Your PDF file must only contain Type 1 or
122 Embedded TrueType fonts. Here are a few instructions to achieve this.

- 123 • You should directly generate PDF files using `pdflatex`.
- 124 • You can check which fonts a PDF files uses. In Acrobat Reader, select the menu
125 Files>Document Properties>Fonts and select Show All Fonts. You can also use the program
126 `pdf fonts` which comes with `xpdf` and is available out-of-the-box on most Linux machines.
- 127 • `xfig` "patterned" shapes are implemented with bitmap fonts. Use "solid" shapes instead.
- 128 • The `\bbold` package almost always uses bitmap fonts. You should use the equivalent AMS
129 Fonts:

130 `\usepackage{amsfonts}`

131 followed by, e.g., `\mathbb{R}`, `\mathbb{N}`, or `\mathbb{C}` for \mathbb{R} , \mathbb{N} or \mathbb{C} . You can also
132 use the following workaround for reals, natural and complex:

```
133 \newcommand{\RR}{\mathbb{R}} %real numbers
134 \newcommand{\Nat}{\mathbb{N}} %natural numbers
135 \newcommand{\CC}{\mathbb{C}} %complex numbers
```

136 Note that `amsfonts` is automatically loaded by the `amssymb` package.

137 If your file contains type 3 fonts or non embedded TrueType fonts, we will ask you to fix it.

138 5.1 Margins in L^AT_EX

139 Most of the margin problems come from figures positioned by hand using `\special` or other
140 commands. We suggest using the command `\includegraphics` from the `graphicx` package.
141 Always specify the figure width as a multiple of the line width as in the example below:

```
142 \usepackage[pdftex]{graphicx} ...
143 \includegraphics[width=0.8\linewidth]{myfile.pdf}
```

144 See Section 4.4 in the graphics bundle documentation ([http://mirrors.ctan.org/macros/](http://mirrors.ctan.org/macros/latex/required/graphics/grfguide.pdf)
145 [latex/required/graphics/grfguide.pdf](http://mirrors.ctan.org/macros/latex/required/graphics/grfguide.pdf))

146 A number of width problems arise when L^AT_EX cannot properly hyphenate a line. Please give LaTeX
147 hyphenation hints using the `\-` command when necessary.

148 References

149 References follow the acknowledgments in the camera-ready paper. Use unnumbered first-level
150 heading for the references. Any choice of citation style is acceptable as long as you are consistent. It
151 is permissible to reduce the font size to `small` (9 point) when listing the references. Note that the
152 Reference section does not count towards the page limit.

153 [1] Alexander, J.A. & Mozer, M.C. (1995) Template-based algorithms for connectionist rule extraction. In
154 G. Tesauro, D.S. Touretzky and T.K. Leen (eds.), *Advances in Neural Information Processing Systems 7*, pp.
155 609–616. Cambridge, MA: MIT Press.

156 [2] Bower, J.M. & Beeman, D. (1995) *The Book of GENESIS: Exploring Realistic Neural Models with the*
157 *GEneral NEural Simulation System*. New York: TELOS/Springer-Verlag.

158 [3] Hasselmo, M.E., Schnell, E. & Barkai, E. (1995) Dynamics of learning and recall at excitatory recurrent
159 synapses and cholinergic modulation in rat hippocampal region CA3. *Journal of Neuroscience* **15**(7):5249-5262.

160 A Appendix / supplemental material

161 Optionally include supplemental material (complete proofs, additional experiments and plots) in
162 appendix. All such materials **SHOULD be included in the main submission**.

163 **NeurIPS Paper Checklist**

164 The checklist is designed to encourage best practices for responsible machine learning research,
165 addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove
166 the checklist: **The papers not including the checklist will be desk rejected.** The checklist should
167 follow the references and follow the (optional) supplemental material. The checklist does NOT count
168 towards the page limit.

169 Please read the checklist guidelines carefully for information on how to answer these questions. For
170 each question in the checklist:

- 171 • You should answer [Yes], [No], or [NA].
- 172 • [NA] means either that the question is Not Applicable for that particular paper or the
173 relevant information is Not Available.
- 174 • Please provide a short (1–2 sentence) justification right after your answer (even for NA).

175 **The checklist answers are an integral part of your paper submission.** They are visible to the
176 reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it
177 (after eventual revisions) with the final version of your paper, and its final version will be published
178 with the paper.

179 The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation.
180 While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a
181 proper justification is given (e.g., "error bars are not reported because it would be too computationally
182 expensive" or "we were unable to find the license for the dataset we used"). In general, answering
183 "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we
184 acknowledge that the true answer is often more nuanced, so please just use your best judgment and
185 write a justification to elaborate. All supporting evidence can appear either in the main paper or the
186 supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification
187 please point to the section(s) where related material for the question can be found.

188 IMPORTANT, please:

- 189 • **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- 190 • **Keep the checklist subsection headings, questions/answers and guidelines below.**
- 191 • **Do not modify the questions and only use the provided macros for your answers.**

192 **1. Claims**

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

195 Answer: **[TODO]**

196 Justification: **[TODO]**

197 Guidelines:

- 198 • The answer NA means that the abstract and introduction do not include the claims
199 made in the paper.
- 200 • The abstract and/or introduction should clearly state the claims made, including the
201 contributions made in the paper and important assumptions and limitations. A No or
202 NA answer to this question will not be perceived well by the reviewers.
- 203 • The claims made should match theoretical and experimental results, and reflect how
204 much the results can be expected to generalize to other settings.
- 205 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
206 are not attained by the paper.

207 **2. Limitations**

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

209 Answer: **[TODO]**

210 Justification: **[TODO]**

211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

- The answer NA means that the paper does not include experiments.

262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315

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

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

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

- 316 • Providing as much information as possible in supplemental material (appended to the
317 paper) is recommended, but including URLs to data and code is permitted.

318 6. Experimental Setting/Details

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

322 Answer: **[TODO]**

323 Justification: **[TODO]**

324 Guidelines:

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

330 7. Experiment Statistical Significance

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

333 Answer: **[TODO]**

334 Justification: **[TODO]**

335 Guidelines:

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

356 8. Experiments Compute Resources

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

360 Answer: **[TODO]**

361 Justification: **[TODO]**

362 Guidelines:

- 363 • The answer NA means that the paper does not include experiments.
- 364 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
365 or cloud provider, including relevant memory and storage.

- 366
- 367
- 368
- 369
- 370
- 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).

371 **9. Code Of Ethics**

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

374 Answer: **[TODO]**

375 Justification: **[TODO]**

376 Guidelines:

- 377
- 378
- 379
- 380
- 381
- 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).

382 **10. Broader Impacts**

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

385 Answer: **[TODO]**

386 Justification: **[TODO]**

387 Guidelines:

- 388
- 389
- 390
- 391
- 392
- 393
- 394
- 395
- 396
- 397
- 398
- 399
- 400
- 401
- 402
- 403
- 404
- 405
- 406
- 407
- 408
- 409
- 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).

410 **11. Safeguards**

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

414 Answer: **[TODO]**

415 Justification: **[TODO]**

416 Guidelines:

- 417
- The answer NA means that the paper poses no such risks.

- 418
- 419
- 420
- 421
- 422
- 423
- 424
- 425
- 426
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
 - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
 - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

427 12. Licenses for existing assets

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

431 Answer: **[TODO]**

432 Justification: **[TODO]**

433 Guidelines:

- 434
- 435
- 436
- 437
- 438
- 439
- 440
- 441
- 442
- 443
- 444
- 445
- 446
- 447
- 448
- The answer NA means that the paper does not use existing assets.
 - The authors should cite the original paper that produced the code package or dataset.
 - The authors should state which version of the asset is used and, if possible, include a URL.
 - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
 - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
 - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

449 13. New Assets

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

452 Answer: **[TODO]**

453 Justification: **[TODO]**

454 Guidelines:

- 455
- 456
- 457
- 458
- 459
- 460
- 461
- 462
- The answer NA means that the paper does not release new assets.
 - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
 - The paper should discuss whether and how consent was obtained from people whose asset is used.
 - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

463 14. Crowdsourcing and Research with Human Subjects

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

467 Answer: **[TODO]**

468 Justification: **[TODO]**

469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496

Guidelines:

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

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

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

Answer: **[TODO]**

Justification: **[TODO]**

Guidelines:

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