

openai / gpt-2

**Join GitHub today**

Dismiss

GitHub is home to over 40 million developers working together to host and review code, manage projects, and build software together.

[Sign up](#)Code for the paper "Language Models are Unsupervised Multitask Learners" <https://openai.com/blog/better-langua...>

50 commits

3 branches

0 packages

0 releases

13 contributors

MIT

Branch: master

[New pull request](#)[Find File](#)[Clone or download](#)

WuTheFWasThat nucleus sampling	Latest commit ac5d522 on Aug 26
src	nucleus sampling last month
.gitattributes	add .gitattributes file to ensure files copied to docker container ha... 7 months ago
.gitignore	updates for 345M model 5 months ago
CONTRIBUTORS.md	Update CONTRIBUTORS.md 7 months ago
DEVELOPERS.md	push 774M model last month
Dockerfile.cpu	push 774M model last month
Dockerfile.gpu	push 774M model last month
LICENSE	add license 8 months ago
README.md	push 774M model last month
domains.txt	add model card last month
download_model.py	push 774M model last month
model_card.md	add model card last month
requirements.txt	Python download script (#89) 7 months ago

**README.md**

**Status:** Archive (code is provided as-is, no updates expected)

## gpt-2

Code from the paper "[Language Models are Unsupervised Multitask Learners](#)".

We have currently released small (124M parameter), medium (355M parameter), and large (774M parameter) versions of GPT-2\*, with only the full model as of yet unreleased. We have also [released a dataset](#) for researchers to study their behaviors.

You can read about GPT-2 and release decisions in our [original blog post](#) and [6 month follow-up post](#).

\* Note that our original parameter counts were wrong due to an error (in our previous blog posts and paper). Thus you may have seen small referred to as 117M and medium referred to as 345M.

## Usage

---

This repository is meant to be a starting point for researchers and engineers to experiment with GPT-2.

For basic information, see our [model card](#).

## Some caveats

- GPT-2 models' robustness and worst case behaviors are not well-understood. As with any machine-learned model, carefully evaluate GPT-2 for your use case, especially if used without fine-tuning or in safety-critical applications where reliability is important.
- The dataset our GPT-2 models were trained on contains many texts with [biases](#) and factual inaccuracies, and thus GPT-2 models are likely to be biased and inaccurate as well.
- To avoid having samples mistaken as human-written, we recommend clearly labeling samples as synthetic before wide dissemination. Our models are often incoherent or inaccurate in subtle ways, which takes more than a quick read for a human to notice.

## Work with us

Please [let us know](#) if you're doing interesting research with or working on applications of GPT-2! We're especially interested in hearing from and potentially working with those who are studying

- Potential malicious use cases and defenses against them (e.g. the detectability of synthetic text)
- The extent of problematic content (e.g. bias) being baked into the models and effective mitigations

## Development

---

See [DEVELOPERS.md](#)

## Contributors

---

See [CONTRIBUTORS.md](#)

## Citation

---

Please use the following bibtex entry:

```
@article{radford2019language,  
  title={Language Models are Unsupervised Multitask Learners},  
  author={Radford, Alec and Wu, Jeff and Child, Rewon and Luan, David and Amodei, Dario and Sutskever, Ilya},  
  year={2019}  
}
```

## Future work

---

We may release code for evaluating the models on various benchmarks.

We are still considering release of the larger models.

## License

---

[MIT](#)