Deep Learning Research Engineer

We are a VC-backed startup creating a system that lets one voice sound like another, focused on very high-quality output and usability for demanding applications like TV and movies, call centers, and gaming. We have projects underway with top Hollywood studios and major call centers. We are graduates of the Techstars startup accelerator.

We are looking for open-minded middle/senior engineers to join us. We really enjoy the stuff we do. We are a team of 9 people. The technical core is based in Kyiv.

Please peruse our job descriptions below, and if you think you could fit into our team do not hesitate to contact us!

Benefits:

- Do real research with deep generative models of audio in a small, scientifically-minded team
- Freedom to try out ideas of your own
- There might be publication opportunities
- Office in the center

Duties:

- Implement and maintain deep learning models in PyTorch
- Train models, keep track of performance, and tune hyperparameters
- Proactively read relevant literature and generate ideas for model enhancements
- Help with deployment and ops as needed

We expect you to:

- Have advanced knowledge in relevant areas of mathematics and statistics
- Have a working knowledge of various techniques, tricks, and best practices for designing deep nets, hyperparameter optimization, visualization, interpretation, etc.
- Have a bunch of failure stories about the deep learning projects you've been involved in (and some success stories)
- Have a good understanding of deep generative models. You may work on autoregressive models, GANs, VAEs, or flow-based models.

Would be a plus:

- Knowledge of Digital Signal Processing theory
- Knowledge of speech processing
- Knowledge of CUDA / interest in developing custom CUDA kernels
- Experience deploying machine learning models
- DevOps experience (Kubernetes, etc.)

Contacts:

Please write to gr+dljob@respeecher.com