

# Joseph Suarez

jsuarez@mit.edu

---

## Massachusetts Institute of Technology | PhD Computer Science *expected 2025*

**Research:** Simulated environments, multiagent intelligence, reinforcement learning. The focus of my current work, Neural MMO, seeks to create neural agents with broad general intelligence by simulating a massively multiagent virtual world

## Stanford University | B.S. Computer Science | GPA: 3.8 / 4.0 2019

**Research:** Two academic years of full time research in natural language processing and computer vision in the labs of Prof. Andrew Ng and Prof. Fei-Fei Li.

## PUBLICATIONS:

**ICML 2020:** *Neural MMO: Ingredients for Massively Multiagent Artificial Open Worlds* (Joseph Suarez, Phillip Isola). Open-source state-of-the-art baseline policies for Neural MMO + infrastructure and utilities.

**AAMAS 2020:** *Neural MMO v1.3: A Massively Multiagent Game Environment for Training and Evaluating Neural Networks* (Joseph Suarez et al.). Structured attentional observation/action processing for efficient training and inference in complex multiagent environments.

**OpenAI Blog and arXiv 2019:** *Neural MMO: A massively multiplayer game environment for intelligent agents* (Joseph Suarez et al.). Framework tested to 100M agent lifetimes. Analysis of agent and species count as magnifiers of environment curriculum and niche formation.

**arXiv 2019:** *GAN you do the GAN GAN?* (Joseph Suarez). Demonstrates that generative adversarial networks can learn distributions over other generative adversarial networks.

**arXiv 2018:** *DDRprog: A CLEVR Differentiable Dynamic Reasoning Programmer* (Joseph Suarez, Justin Johnson, Li Fei-Fei). State-of-the-art on CLEVR, differentiable hard attention and forking techniques for dynamic assembly and execution of neural program trees.

**NIPS 2017:** *Language Modeling with Recurrent Highway Hypernetworks* (Joseph Suarez). State-of-the-art on Penn Treebank, detailed analysis of gradient flow in recurrent architectures.

**arXiv 2017:** *Efficient Approaches to Batch Parallelization for Dynamic Neural Architectures* (Joseph Suarez, Clare Zhu). 10x speedup on state-of-the-art CLEVR visual question answering architecture. 1000x speedup on sparsely gated mixture of experts layers.

## ADDITIONAL EXPERIENCE:

**Summer 2018:** 6 month full time machine learning research internship at **OpenAI**. Developed Neural MMO, a massively multiagent platform for training reinforcement learning agents.

**Summer 2017:** Worked with **Justin Johnson** in the **Stanford Vision lab** on differentiable and dynamic approaches to CLEVR Visual Question Answering.

**Summer 2016:** Intern at **LucidCam**. Prototyped augmented reality demos, assisted with various computer vision based tasks.