Mikhail Jacob

bitbucket.org/mikkujacob | github.com/mikhailjacob
mikhailjacob.in | mikkujacob@gmail.com
+46 722999253 | Stockholm, Sweden

Lead ML engineer + human-centred AI researcher

Education

PhD Computer Science, Georgia Institute of Technology, December 2019
PhD Dissertation: "Improvisational Artificial Intelligence for Embodied Co-creativity"
MS Computer Science, Georgia Institute of Technology, May 2013, M.S. Project: "Viewpoints AI"
BE Computer Science & Engineering, Manipal Institute of Technology, July 2011

Skills & Competencies

R&D: machine learning research, system design, engineering, evaluation, & deployment | writing & presentation
ML Competencies: deep reinforcement learning (RL) | imitation learning | generative models | deep learning
Frameworks: PyTorch | TensorFlow | Keras | Ray/RLlib | Stable Baselines 3
Languages: Python | C# | C++ | Java | JavaScript
Tools: Unity | Unreal Engine | Processing & P5.js | GCP

Professional Experience

Lead Machine Learning Engineer | Machine Learning Team | **Resolution Games**, Stockholm, Sweden | September 2021 – Present **| Machine Learning in XR Games**

- Summary: Started and led the ML team to develop game features for 1 AR and 2 VR games + 2 prototypes.
- Two years delivering ML bots for VR games using deep RL and imitation learning (RLlib + Unity) on device.
- Shipped passthrough AR interactive dancing NPC experience on Quest 2 using on-device deep generative models (PyTorch + CVAEs) to generate real-time character animation and appoximate full-body IK from VR sensors.
- Developed 2 game prototypes integrating large language models and latent diffusion models for speech to image generation and speech-controlled NPC plan generation.
- Developed entire ML strategy and roadmaps for Resolution with senior/executive stakeholders + led recruitment.
- Mentored and supervised junior ML team colleagues to grow their competence and independence.
- Effectively shipped features with cross-functional teams including game developers, artists, producers, executives.

Researcher | Deep Reinforcement Learning for Games | **Microsoft Research**, Cambridge, UK | August 2019 – August 2021 | **Reinforcement Learning for Game Designers** | aka.ms/RLProblems4Gaming

- Summary: End to end ML research on designer reinforcement and imitation learning by studying workflows + challenges; identifying opportunities for impact; and redefining RL workflows through prototype models + tools.
- Studied 17 game AI/ML creators using interviews + thematic analysis to identify 10 challenges and opportunities for researchers in commercial games with Best Paper Award at AIIDE 2020 conference.
- Implemented/trained InfoGAIL (inverse RL with GANs) to adapt model w/o retraining by interpolating experts.
- Mentored PhD research intern on adapting agent aesthetic style using preference learning + potential-based reward shaping (PBRS) + automatic reward weighting with a Microsoft Research blog post.
- Redesigned game designer workflows with RL, created human-in-the-loop RL algorithms, and made prototype designer tool to improve model robustness from domain distribution shift for product partner game scenario.
- Redesigned game designer workflows with RL, created multi-task RL algorithms, and made prototype designer tool in Unreal Engine to reduce game designer behaviour exploration time for product partner game scenario.
- AIIDE 2020 best paper award, 2 publications, 1 blog post, 3 prototype tools, 3 invited talks, 1 PhD intern.

Graduate Research Assistant | Expressive Machinery Lab | **Georgia Institute of Technology**, Atlanta, USA | May 2012 – July 2019

Research: The Robot Improv Circus | tiny.cc/RobotImprovCircus | August 2016 – July 2019

- Summary: Research team lead for human-AI improv theatre VR installation to play 'Props' (miming pretend actions with abstract props as real objects) using *creative arc negotiation* to quickly choose actions from many options.
- Completed physical, virtual environment, and interaction design of the Robot Improv Circus installation (Unity).

- Collected human data, designed deep generative model architecture (conditional VAE) to generate action variants, and trained + evaluated models on mimed actions + prop attributes.
- Architect + engineering lead for ML system performing creative arc negotiation using action generation (CVAE model inference), improvisational strategies inspired by human improvisers (search strategies in CVAE latent space), and heuristic models for evaluating novelty (RTree semantic distance to comparable experiences), surprise (Bayesian surprise + expectation deviation), and value (smoothness + recognisability) of potential action choices.
- Evaluation studies showed that both observers and participants strongly preferred creative arc negotiation for agent creativity and logical coherence, while observers also enjoyed it more, compared to our baseline.
- C&C 2021 best paper award, Creative Curricular Initiatives (CCI) $6,000 grant awarded, 2 invited installations, 4 publications, 2 invited talks.

Research: LuminAI (formerly Viewpoints AI) | tiny.cc/LuminAI | August 2012 – August 2017

- Summary: Research team lead for human-AI improvised dance installation with agent learning novel dance moves and patterns of interaction from people.
- Completed physical, character, and interaction design of LuminAI installations (Java + Processing, later Unity).
- Architect + engineering lead for agent architecture to improvise dance using movement clustering (for generalisation), fluid turn-taking (combining rhythm + pause + turn classification), case-based imitation learning (memorise, imitate, reuse movement clusters + patterns in future), and human improvisation reasoning strategies.
- ACCelerate Festival 2017 selection to the Smithsonian Institution National Museum of American History, Field Experiment ATL grant + showcase finalist, Neukom Institute Turing Test in Creative Arts 2017: DanceX Prize winner, 11 national + international invited installations, 7 peer-reviewed publications, and 4 invited talks.

Research: Computational Representations of Play | tiny.cc/ComputationalPlay | August 2012 – August 2016

- Summary: Created a cognitive model of conceptual blending to enable human-AI gestural pretend play with toys.
- Created a system for blending semantic knowledge as a model of human-robot (or virtual agent) pretend toy play.
- 3 peer-reviewed publications, 1 prototype system, and 2 invited talks.

Additional Projects

Player Assist Bot | *Microsoft Global Hackathon 2020* | tiny.cc/PlayerAssistBotTraining | Summer 2020

- Team + engineering lead for end to end creation of RL-trained bot system for assisting players on request.
- Used game video capture + controller emulation for game integration without developer API (compatibility) and TensorFlow convolutional neural network (CNN) models for reward signal classification from game video.
- Trained and demonstrated (RLLib) bot on Streets of Rage 4. Project awarded Honourable Mention globally.

Flame Warz – A Twitter Conflict Game | *Advanced Game AI* | Fall 2012

- Created a game with generated quests; characters + items mined from Twitter trends (people, followers, items).
- Created villain's emotion model combining plan appraisals (EMA) and PAD space moods (ALMA) that controlled the game ending and allowed multiple conflict resolution styles like diplomacy, aggression, shaming, etc.

Generative Abstract Art Game | *Game AI* | Spring 2012

- Created a procedurally generated game (Unity) that created abstract art levels using a genetic algorithm.
- Opponent parameters and behaviour were adapted dynamically using simple player modelling.

RadVenture | *AI Storytelling In Virtual Worlds* | Fall 2011

- Created an alternate reality game engine (Android) with quest generation + management using HTN planning.

Selected Publications*

- **Jacob, M**., & Magerko, B. (2021). "Empirically Evaluating Creative Arc Negotiation for Improvisational Decision-Making." Creativity and Cognition. Virtual Event, Italy. C&C 2021 Best Paper Award.
- **Jacob, M.**, Devlin, S., & Hofmann, K. (2020). "'It's Unwieldy and It Takes a Lot of Time' - Challenges and Opportunities for Creating Agents in Commercial Games." In the Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE) 2020. Alberta, Canada. AIIDE 2020 Best Paper Award.
- **Jacob, M.** (2019). "Improvisational artificial intelligence for embodied co-creativity." Doctoral dissertation, Georgia Institute of Technology 2019. Atlanta, USA.
- **Jacob, M.**, Chawla, P., Douglas, L., He, Z., Lee, J., Sawant, T., & Magerko, B. (2019). "Affordance-based generation of pretend object interaction variants for human-computer improvisational theater." In the Proceedings of the 10th ACC International Conference on Computational Creativity (ICCC) 2019, Charlotte, USA.

* Full publication list available here.