Dr. Dan A. Calian

dan.andrei.calian@gmail.com
Personal Website
Google Scholar Profile

Education

2012
-
2018
University College London, Department of Computer Science
Doctor of Engineering (PhD equiv.)
Thesis: "Customised Light Probes and Inverse Lighting Methods for Relighting"   PDF

My thesis introduces novel inverse lighting methods for estimating the lighting conditions of a real scene given a photograph of a known object as input. Two physics-based methods based on custom light probes ("shading probes") are proposed. A deep-learning approach leveraging priors on human faces is introduced. A novel formulation for fast convolutional sparse coding on the sphere (leveraging state-of-the-art spherical harmonics theory) is also introduced.

2011
-
2012
University College London
Master of Research in Virtual Environments, Imaging & Visualisation - Merit (75%)
Dissertation on "Context Aware Rendering". Group project accepted to SIGGRAPH Asia.

Key modules: Computational Photography and Capture (84), Advanced Modelling, Rendering and Animation (92), Computer Graphics (93).

2008
-
2011
University of Nottingham
BSc (Hons) Computer Science with Artificial Intelligence - First Class (81%)
Graduated first in my year. Dissertation published as journal article.

Key modules: Machine Learning (80), Artificial Intelligence Methods (96) & Programming (86), C/C++ for Java Programmers (95), Algorithms and Data Structures (89).

Work Experience

Since
09/2019
Research Engineer - DeepMind, London

- Co-led two research projects on constrained reinforcement learning, with two joint first-authored paper's submitted. In the first project we investigated using metagradients for tuning Lagrange multipliers. In the second project we proposed an algorithm for training constraint-satisfying policies to be robust to environment perturbations.

- Contributed to a research project on neural network verification, resulting in a KDD 2020 publication.

- My responsibilities include: discussing and proposing research ideas; initiating collaborations; implementing algorithms; setting up infrastructure and baselines; running and analysing experiments; giving internal research presentations and writing academic papers.

2018 - 2019	**Research Scientist** - Blue Prism, Research Lab, London • Led applied research on integrating state-of-the-art object detection into company's core product; at its busiest, the project consisted of five team members. Results delivered as on-stage demo at the company's public conference "Blue Prism World" (ExCel London). • Co-led research on spatially-coherent randomised attention maps. • Helped set the group's research agenda. • Managed a junior team member (mentoring, feedback, performance). • Mentored several junior team members. • Conducted telephone and on-site interviews; set technical interview questions.
08/2017 - 12/2017	**Research Intern** - Disney Research, Los Angeles • Derived a novel theoretical formulation for frequency-domain convolutional sparse coding for spherical signals. • Adapted ℓ_0-constrained least squares algorithms to validate formulation: compressive sampling matching pursuit, hard-thresholding pursuit and gradient descent with sparsification (Python: Theano, NumPy and C++: Eigen).
2016 - 2017	**Machine Learning Engineer** - Onfido, London • Implemented PatchMatch for copy-paste fraud detection. Wrote highly-optimised C++ which could process 20MP images in under 2 minutes (in production). • Trained convolutional neural networks on fraud-heatmap outputs (TF). • Analysed encountered fraud types using Bernoulli Mixture Models, and shared findings. • Presented papers at research reading groups; presented research findings in a company-wide meeting; participated in daily scrums; liaised with product managers. *I consulted initially, then worked part-time, and finally full-time for three months.*
04/2016 - 07/2016	**Research Intern** - Disney Research, Pittsburgh • Led research on estimating non-negative lighting from *outdoor* face photographs. • Parametrised lighting as latent space of convolutional autoencoder trained on skies (Theano) and through a physically-based parametric model. • Ported and used Mark Schmidt's L-BFGS solvers (`minFunc`, `minConf`) from MATLAB to Python (NumPy). • Implemented convolutional neural networks for spherical signals in the spatial domain, with sparse precomputed dot-product matrices (Theano).

07/2014
-
10/2014

Research Intern - Disney Research, Pittsburgh
- Led research on estimating non-negative lighting from *indoor* face photographs using: spherical and hemispherical harmonics, Haar wavelets and Voronoi wavelets.
- Implemented proximal gradient methods for alternating albedo and lighting optimisations (MATLAB, Mosek, C++).
- Encoded priors on skin reflectance as Gaussian Mixture Models and used them to constrain optimisations.
- Used closures in MATLAB to implement automatic differentiation (chain-rule) for using iterative solvers (L-BFGS).
- Implemented facial geometry estimation by deforming a template mesh onto 3D facial landmarks (MATLAB, dlib).
- Implemented real-time inverse lighting from arbitrary 3D printed objects for internal Disney Research demo (C# and C++: Unity Engine, Ceres, Eigen).

08/2012
-
09/2012

Research Intern - Disney Research, Edinburgh
- Worked on augmented reality-related projects: inverse lighting from voxel objects and ARDrone control (C#, Unity).
- Work on optical illumination computation and capture patented by Disney.

06/2010
-
08/2010

Research Assistant - School of Computer Science, University of Nottingham
- Implemented "memetic" extensions to a genetics-based learning classifier system (C++).
- Internship served as a starting point for my BSc dissertation.

Publications & Working Papers

- **Dan Andrei Calian**, Florian Stimberg, Olivia Wiles, Sylvestre-Alvise Rebuffi, Andras Gyorgy, Timothy Mann, Sven Gowal. *Defending Against Image Corruptions Through Adversarial Augmentations.* arXiv preprint, Apr 2021. PDF

- Sylvestre-Alvise Rebuffi, Sven Gowal, **Dan Andrei Calian**, Florian Stimberg, Olivia Wiles, Timothy Mann. *Fixing Data Augmentation to Improve Adversarial Robustness.* arXiv preprint, March 2021. PDF

- Daniel Mankowitz*, **Dan Andrei Calian***, Rae Jeong, Cosmin Paduraru, Nicolas Heess, Sumanth Dathathri, Martin Riedmiller, Timothy Mann. *Robust Constrained Reinforcement Learning for Continuous Control with Model Misspecification.* arXiv preprint, Oct 2020. PDF

- **Dan Andrei Calian***, Daniel Mankowitz*, Tom Zahavy, Zhongwen Xu, Junhyuk Oh, Nir Levine, Timothy Mann. *Balancing Constraints and Rewards with Meta-Gradient D4PG.* arXiv preprint, Oct 2020 (to be published in ICLR proceedings). PDF

- Anton Zhernov, Krishnamurthy Dj Dvijotham, Ivan Lobov, **Dan Andrei Calian**, Michelle Gong, Natarajan Chandrashekar, Timothy Mann. *The NodeHopper: Enabling Low Latency Ranking with Constraints via a Fast Dual Solver* Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2020/8/23. PDF

- **Dan Andrei Calian***, Peter Roelants*, Jacques Calì, Ben Carr, Krishna Dubba, John E. Reid, Dell Zhang. *SCRAM: Spatially Coherent Randomized Attention Maps.* arXiv preprint, May 2019. PDF

- **Dan Andrei Calian** et al. *Fast Convolutional Sparse Coding on the Spherical Domain.* *Unpublished, work-in-progress.* 2019. PDF

- **Dan Andrei Calian**, Jean-François Lalonde, Paulo Gotardo, Tomas Simon, Iain Matthews, Kenny Mitchell. *From Faces to Outdoor Light Probes.* Computer Graphics Forum (Proc. Eurographics), 37(2):51-61, 2018. PDF

- **Dan Andrei Calian**, Kenny Mitchell, Derek Nowrouzezahrai, Jan Kautz. *The Shading Probe: Fast Appearance Acquisition for Mobile AR.* Proceedings of SIGGRAPH Asia 2013 Technical Briefs. ACM, New York, NY, USA, Article 20. PDF

- Jacques Calì, **Dan Andrei Calian**, Cristina Amati, Rebecca Kleinberger, Anthony Steed, Jan Kautz, and Tim Weyrich. *3D-Printing of non-assembly, articulated models.* ACM Transactions on Graphics (Proc. SIGGRAPH Asia), 31(6):130:1–130:8, 2012. PDF

- **Dan Andrei Calian** and Jaume Bacardit. *Integrating memetic search into the BioHEL evolutionary learning system for large-scale datasets.* Memetic Computing. In press, 2013. PDF

Academic Reviewing

I have been a reviewer or a member of the programme committee for:

- Robust and Reliable Machine Learning in the Real World Workshop at ICLR (**PC**: 2021)

- Challenges of Real World Reinforcement Learning Workshop at NeurIPS (2020)

- Intelligent Process Automation Workshop at AAAI (2020)

- i3D - ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (**PC**: 2019 & 2018)

- Eurographics - The Annual Conference of the European Association for Computer Graphics (2019)

- ISMAR - IEEE International Symposium on Mixed and Augmented Reality (2016)

- TVCG - IEEE Transactions on Visualization and Computer Graphics (2014)

Awards

- HackerRank "Machine Learning Codesprint" on Natural Language Processing: 1^{st} place (2017)

- UCL Graduate School Student Conference Fund (2012)

- EPSRC Engineering Doctorate Funding (2011)

- EPSRC Vacation Bursary (2010)

- Student Course Representative Bronze Award 2008/2009 and 2009/2010

- C/C++ and Oracle Certificate of Professional Competency (2008)

- TOEFL iBT 108/120 (2007)

- SAT Subject Tests: Mathematics II 760/800 (2007)

Technical Skills

- Excellent programming skills in a large number of languages: C++, Python, MATLAB, C#, bash shell, GLSL/Cg, Java, etc.

- Excellent mathematical modelling and optimisation knowledge as well as practical knowledge in using related frameworks & libraries for implementation: (Python) Jax, TF, PyTorch, Theano, NumPy, SciPy; (C++) Eigen, Ceres; (MATLAB) optimisation toolbox, etc.

- During my undergraduate degree I wrote several 3D and 2D games to improve my understanding of computer graphics and A.I.: Rubik's Cube 3D (C++), Backgammon 3D with A.I. (C#), Circling Rabbit (C++) and Pong with A.I. (Java).

My personal interests include: hiking, skiing, playing guitar and travelling.