Hsiao-Yu Fish Tung

Carnegie Mellon University, Machine Learning Department

5000 Forbes Ave., Pittsburgh, PA, USA
+1 (412) 9611335
htung@cs.cmu.edu
http://sfish0101.bitbucket.io/

Research Interests

Deep Learning for Computer Vision, Graphics and Language.

Current Research Projects

Self-supervised Learning of Motion Capture. (NIPS2017, Spotlight) [9]

We propose a learning based, end-to-end motion capture model for monocular videos in the wild. Current state of the art solutions for motion capture from a single camera are optimization driven, which are susceptible to local minima. This has been the bottleneck that forced using clean green-screen like backgrounds at capture time, manual initialization, or switching to multiple cameras as input resource. Instead of optimizing mesh and skeleton parameters directly, our model optimizes neural network weights that predict 3D shape and skeleton configurations given a monocular RGB video. Our model is trained using a combination of strong supervision from synthetic data, and self-supervision from differentiable rendering of (a) skeletal keypoints, (b) dense 3D mesh motion, and (c) human-background segmentation, in an end-to-end trainable framework. Self-supervision by back-propagating through differentiable rendering allows (unsupervised) adaptation of the model to the test data, and offers much tighter fit that a pretrained fixed model. We show that the proposed model improves with experience and converges to low error solutions where previous optimization methods fail.

Adversarial Inverse Graphics Networks: Learning 2D-to-3D Lifting andImage-to-Image Translation from Unpaired Supervision. (ICCV2017, Project Website: https://sites.google.com/view/adversarial-inversion) [8]

We propose adversarial inversion graphics network, a weakly supervised neural network model that combines self-supervision with adversarial constraints. Given visual input, our model first generates a set of desirable intermediate latent variables, which we call "imaginations" (e.g., 3D pose and camera viewpoint). Then a differentiable renderer projects these imaginations to reconstruct the input, and discriminator networks constrain the imaginations, using corresponding reference repositories, to reside in the right "domain." Our model is trained to minimize reconstruction and adversarial losses and can be trained without paired annotations. We empirically show adversarial inversion outperforms previous state-of-the-art supervised models on 3D human pose estimation from a static image, and 3D scene depth estimation from per-frame motion. Further, we show results on biased image editing.

Generative Models and Model Criticism via Optimized Maximum Mean Discrepancy. (ICLR2017) [4]

We propose a method to optimize the representation and distinguishability of samples from two probability distributions, by maximizing the estimated power of a statistical test based on the maximum mean discrepancy (MMD). This optimized MMD is applied to the setting of unsupervised learning by generative adversarial networks (GAN), in which a model attempts to generate realistic samples, and a discriminator attempts to tell these apart from data samples. In this context, the MMD may be used (1) as a discriminator for distribution matching and (2) it can be used to evaluate the performance of a generative model, by testing the model's samples against a reference data set. Arxiv: https://arxiv.org/abs/1611.04488

Education

2015 - now	**Carnegie Mellon University** PhD in Machine Learning (QPA:4.03) Advisor: Prof. Katerina Fragkiadaki & Prof. Alexander J. Smola
2013 – 2015	**Carnegie Mellon University** M.S. in Machine Learning (QPA:4.0) Advisor: Prof. Alexander J. Smola
2009-2013	**National Taiwan University** B.S. in Electrical Engineering Ave. GPA: 3.95/4.0; Rank: 9%

Working Experience

2017 summer	**Adobe Research** Advisor: Ersin Yumer
2016 summer	**Google Brain Team, Google Inc.** Advisor: Andrew Dai
2015 summer	**Parallel Computing Lab, Intel Labs** Advisor: Shang Li
2014 summer	**Research Assistant, CMU** Advisor: Aarti Singh
2012 summer	**Software Engineer, MediaTek** Home Entertainment group
2011-2013	**Research Assistant, NTU** Advisor: Soo-Chang Pei, Tian-Li Yu

Awards

2017	NIPS Travel Award
2014	NIPS Travel Award
2013	KDD Cup Award, Track1& 2 Champion
2013	NTU Student Outstanding Performance Scholarship
2012	NTU President Award, 1/245

Coursework & Teaching

Teaching Assistant at CMU:

10-715	Advanced Introduction to Machine Learning Machine Learning course for ML PhD students
10-601	Introduction to Machine Learning Machine Learning course for master students

CMU selected courses:

Visual Learning and Recognition,
Statistical Machine Learning,
Intermediate Statistics, Multimedia Database,
Advanced Optimization and Randomized Methods,
Foundations of Machine Learning and Data Science

Skills

Deep-learning	Tensorflow • Theano • Caffe • Torch
Programming	Python • C++ • Matlab • Shell
Language	Chinese(native), English (fluent)

Past Projects

Generative Adversarial Nets (GANs) for Text.
Project with Google Brain
GANs has been applied successfully on images generation, however, it is still unclear how to apply it on sequential and discrete data like text. In this project, we explored a wide range of structures and training strategies for sentence generation using GANs. We concluded that, for sentence generation, it is not practical to apply GANs directly on the data space because the training process becomes unstable and slow. Instead, we designed a model that uses moment matching method to constraint on the continuous coding space of an autoencoder for sentences. Our model is able to generate reasonable sentences from the model. Besides, we showed that, given two sentences, the model is able to generate a novel sentence that is an interpolation between the given sentences.

Accelerating Long Short Term Memory Network.
Project with Intel Parallel Computing Labs
We developed a fast and scalable LSTM implementation on Intel latest CPU Platform. Our CPU-based system has 4 times higher efficiency than GPU based systems. We further show hierarchical lstm has more potential in parallelism and can achieve better performance on machine translation, document classification and question−answering.

Spectral Methods for Indian Buffet Process (IBP) and Hierarchical Dirichlet Processes (HDP). (NIPS 2014, NIPS2015 workshop, preprint journal on Arxiv.) [6] [7] [10]
Proposed efficient spectral algorithm for non-parametric model inference. The algorithm provides superior accuracy and cheaper computation than comparable Variational Bayesian approach on a number of problems with reconstruction guarantees. We further generalized the tensors to multi-layer HDP. The algorithm is applied on large dataset: Reviews, Enron Email and proved to have better performance than standard spectral LDA.

Fast and Guaranteed Tensor Decomposition via Sketching. (NIPS 2015) [11]
Proposed fast and randomized tensor CP decomposition algorithms based on sketching. The idea is to use count sketches for randomized computation of tensor contractions via FFTs, without explicitly forming the tensors. The proposed method speed up the naive implantation by a factor of 10. Besides, different from general sampling method, the estimator obtained using the proposed method is supported by theoretical bounds.

2013 KDD Cup Award, Track1&2 Champion. [1] [3] [2]
Algorithm@National Taiwan University, Advisor: Chih-Jen Lin
Large-scale datasets which includes 250K authors and 2.5M papers from Microsoft Academic Search is used in the competition. Track1 aims at distinguishing whether the assigned papers are truly written by a given author with some training data. Track2 focuses on determining duplicated authors which is a unsupervised learning task. The method successfully make 30-40x speeds up of the power tensor method.

Novel Traffic Light Timing Adjustment Strategy Based On Genetic Algorithm.(CEC 2013, Oral) [5]
Solved traffic signal timing optimization problem which aims at shortening the average traffic time using linear regression to learn from the global optimal solution obtained by GA. The new adjustment strategy outperforms recent optimization methods.

References

[1] W.-S. Chin, Y.-C. Juan, Y.-Zhuang, H.-Y. Tung Felix Wu, T. Yu, J.-P. Wang, C.-X. Chang, C.-P. Yang, W.-C. Chang, K.-H. Huang, T.-M. Kuo, S.-W. Lin, Y.-S. Lin, Y.-C. Lu, Y.-C. Su, C.-K. Wei, T.-C. Yin, C.-L. Li, T.-W. Lin, C.-H. Tsai, S. d. Lin, H.-T. Lin, and C.-J. Lin. Effective string processing and matching for author disambiguation. *KDD CUP 2013 workshop, KDD*, 2013.

[2] W.-S. Chin, Y.-C. Juan, Y.-Zhuang, H.-Y. Tung Felix Wu, T. Yu, J.-P. Wang, C.-X. Chang, C.-P. Yang, W.-C. Chang, K.-H. Huang, T.-M. Kuo, S.-W. Lin, Y.-S. Lin, Y.-C. Lu, Y.-C. Su, C.-K. Wei, T.-C. Yin, C.-L. Li, T.-W. Lin, C.-H. Tsai, S. d. Lin, H.-T. Lin, and C.-J. Lin. Effective string processing and matching for author disambiguation. *Journal of Machine Learning Research*, 2014.

[3] C.-L. Li, Y.-C. Su, T.-W. Lin, C.-H. Tsai, W.-C. Chang, K.-H. Huang, T.-M. Kuo, S.-W. Lin, Y.-S. Lin, Y.-C. Lu, C.-P. Yang, C.-X. Chang, W.-S. Chin, Y.-C. Juan, H.-Y. Tung, J.-P. Wang, C.-K. Wei, Felix Wu, T.-C. Yin, T. Yu, Y. Zhuang, S. d. Lin, H.-T. Lin, and C.-J. Lin. Combination of feature engineering and ranking models for paper-author identification in kdd cup 2013. *KDD CUP 2013 workshop, KDD*, 2013.

[4] Dougal J. Sutherland, Hsiao-Yu Tung, Heiko Strathmann, Soumyajit De, Aaditya Ramdas, Alex Smola, and Arthur Gretton. Generative models and model criticism via optimized maximum mean discrepancy. *ICLR*, 2017.

[5] Hsiao-Yu Tung, Wei-Chui Ma, and Tian-Li Yu. Novel traffic light timing adjustment strategy based on genetic algorithm. *IEEE Congress on Evolutionary Computation*, 2014.

[6] Hsiao-Yu F. Tung and Alexander J. Smola. Spectral methods for indian buffet process inference. *NIPS*, 2014.

[7] Hsiao-Yu F. Tung, Chao-Yuan Wu, Manzil Zaheer, and Alexander J. Smola. Spectral methods for the hierarchial dirichlet process. *NIPS 2015 Workshop on Large Scale Representation Learning*, 2015.

[8] Hsiao-Yu Fish Tung, Adam Harley, William Seto, and Katerina Fragkiadaki. Adversarial inverse graphics networks: Learning 2d-to-3d lifting andimage-to-image translation from unpaired supervision. *ICCV*, 2017.

[9] Hsiao-Yu Fish Tung, Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. *NIPS*, 2017.

[10] Hsiao-Yu Fish Tung, Chao-Yuan Wu, Manzil Zaheer, and Alexander J. Smola. Spectral methods for nonparametric models. 2017.

[11] Yining Wang, Hsiao-Yu Tung, Alexander Smola, and Animashree Anandkumar. Fast and guaranteed tensor decomposition via sketching. *NIPS*, 2015.