

Intel Collaborative Research Institute for Computational Intelligence

Deep Learning + Visual Day

Ronny Ronen
May 9, 2017

ICRI-CI web-site: <http://icri-ci.technion.ac.il/>



ICRI-CI Research Projects (5/2017)

Conversational Understanding 8

1. Universal Semantics (UCCA) **Ari Rappoport (H)**
2. Automatic Measurement of Transcription Quality **Moshe Koppel (B)**
3. Holistic Inference for Natural Language Processing Amir **Globerson (H)**
4. Open Information Extraction Knowledge Graphs **Ido Dagan (B)**
5. Unsupervised Extraction of Relations and Events **Ronen Feldman (H)**
6. Hybrid Models for Minimally Supervised Information Extraction from Conversations **Roi Reichart (T)**
7. Syntactic and Semantic Reranking of Speech Interaction Data **Yoav Goldberg (B)**
8. Topic Dependent Language Modeling **Jacob Goldberger (B)**

Moshe Wasserblat

Optimized Intel Architecture 8

1. Memory Traffic Reduction for Big Data & ML **Uri Weiser (T)**
2. Accelerators for Big Data & Machine Learning **Ran Ginosar (T), Oded Schwartz (H)**
3. Memory Intensive Architectures **Shahar Kvatinsky (T), Y. Cassuto, E. Friedman (T)**
4. Context-based Prefetching using Reinforcement Learning **Yoav Etsion (T)**
5. Software managed paging in secured memory **M. Silberstein (T)**

Debbie Marr

Optimized Deep Learning 11

1. Optimal Deep Learning & the Information Bottleneck Principle **Naftali Tishby (H)**
2. SimNets: A Generalization of Convolutional Networks **Amnon Shashua (H)**
3. Rigorous Algorithms for Distributed Deep Learning **Shai Shalev Shwartz (H)**
4. Mega-Class Efficient Deep Learning **Koby Cramer (T)**
5. Outlier Robust Distributed Learning **Shie Mannor (T)**
6. Unsupervised and Semi-supervised Ensemble Learning **Boaz Nadler (W)**
7. Distributed Methods for Non-Convex and Deep Learning **Ohad Shamir (W), Nati Srebro (T)**
8. Scene Understanding: from Image to Text and from Image & Question to an Answer; Efficient Optical Flow **Lior Wolf (TA)**
9. Image Restoration w/ Deep Learning **M. Zibulevsky (T)**
10. Using Deep Learning to Medical Imaging **Hayit Greenspan (TA)**

Ronny Ronen

Visual Systems / Other 4

1. Blind Video: Video Without Photographers **Shmuel Peleg (H)**
2. Mental Phenotyping from 3D Cameras **Daphna Weinshall (H)**
3. Statistics of Depth Images **Yair Weiss (H)**
4. Providing People with Arguments during Persuasive Discussion **Sarit Kraus (B)**

Ronny Ronen



ICRI-CI
Intel Collaborative Research Institute
Computational Intelligence



Deep Learning Research

- **Technology to leapfrog Deep Learning**

- Core DL technology: SimNets, DeepRL¹, RNN², ...
- Distributed Learning: non-Convex, Geometric Structure, ...
- Advanced DL application: Scene understanding, Optical flow, Medical imaging, Restoration

- **Understanding Deep Learning**

- Suggested a focus topic last year by ICRI-CI BOA
- About to publish a “paper bundle” consisting of 5 papers, Titled:

***Why & When Deep Learning Works:
Looking Inside Deep Learning***

- Bundle URL: <https://arxiv.org/abs/1705.03921>

¹ DeepRL: Deep Reinforcement Learning

² RNN: Recurrent Neural Network

Research Activities

Core
Technology

Distributed
Algorithms

Advanced
DL Apps.

| Researcher | Univ. | Research Title |
|--------------------------------|-----------------|--|
| Naftali Tishby | Hebrew U | Opening the black box of Deep Neural Networks via Information |
| Amnon Shashua | Hebrew U | SimNets: A Generalization of Conv Nets Deeper Understanding of Deep Learning |
| Boaz Nadler | Weizmann | Un/Semi-Supervised Ensemble Learning |
| Shie Mannor | Technion | Outlier robust distributed learning; Deep Q Network |
| Shai Shalev-Shwartz | Hebrew U | Novel Deep Learning algorithms |
| Ohad Shamir & Nathan Srebro | Weizmann TTI | Distributed Methods for non-Convex and Deep Learning Exploiting Geometric Structure |
| Koby Crammer | Technion | Mega class Deep Learning |
| Lior Wolf | Tel Aviv U | Scene understanding: from image to text and question to an answer; Efficient optical flow |
| Hayit Greenspan | Tel Aviv U | Deep Learning for Medical Imaging |
| Mark Zibulevsky | Technion | Image restoration using Deep Learning |

Deep Learning Research

- **Technology to leapfrog Deep Learning**
 - Core DL technology: SimNets, DeepRL¹, RNN², ...
 - Distributed Learning: non-Convex, Geometric Structure, ...
 - Advanced DL application: Scene understanding, Optical flow, Medical imaging, Restoration
- **Understanding Deep Learning**
 - Suggested as a focus topic last year by ICRI-CI BOA
 - About to publish a “paper bundle” consisting of 5 papers, Titled:

Why & When Deep Learning Works: Looking Inside Deep Learning

- Bundle URL: <https://arxiv.org/abs/1705.03921>

¹ DeepRL: Deep Reinforcement Learning

² RNN: Recurrent Neural Network

“Why & When Deep Learning Works: Looking Inside Deep Learning”

- Available at <https://arxiv.org/abs/1705.03921>

| Authors | Title |
|---|---|
| <i>Understating of why and when deep networks work (and do not work):</i> | |
| Naftali Tishby, et.al. | <u>Opening the Black Box of Deep Neural Networks via Information</u> |
| Shai Shalev-Shwartz, Ohad Shamir, et.al. | <u>Failures of Gradient-Based Deep Learning</u> |
| <i>The impact of geometry on the expressiveness of deep networks:</i> | |
| Amnon Shashua, Nadav Cohen, et.al. | <u>Analysis and Design of Convolutional Networks via Hierarchical Tensor Decompositions</u> |
| Nathan Srebro, et.al. | <u>Geometry of Optimization and Implicit Regularization in Deep Learning</u> |
| <i>Interpretability of deep networks:</i> | |
| Shie Mannor, et.al. | <u>Graying the black box: Understanding DQNs</u> |

Visual Systems / Other 4

1. Blind Video: Video Without Photographers
Shmuel Peleg (H)
2. Mental Phenotyping from 3D Cameras
Daphna Weinshall (H)
3. Statistics of Depth Images
Yair Weiss (H)
4. Providing People with Arguments during Persuasive Discussion
Sarit Kraus (B)

Today's Talks

Deep Learning I

| | |
|----------------|---|
| Naftali Tishby | Opening the black box of Deep Neural Networks via Information: A deeper theory and some new algorithms |
| Amnon Shashua | Expressive efficiency and inductive bias of convolutional networks: the use of hierarchical tensor decompositions for network design and analysis |
| Ohad Shamir | Failures of Gradient-Based Deep Learning |
| Shie Mannor | End-to-end Deep Imitation Learning |

Deep Learning II + Visual

| | |
|---------------------|---|
| Lior Wolf | Stereo Matching, Optical Flow, and Filling the Gaps |
| Shai Shalev-Shwartz | Deep reinforcement learning for driving policy |
| Hayit Greenspan | Deep Learning in Medical Imaging: The Data Challenge |
| Daphna Weinshall | Implicit Media Tagging and Affect Prediction from RGB-D video of spontaneous facial expressions |
| Yair Weiss | The Return of the Gating Network: Combining Discriminative and Generative Training in models of RGBD Images |