WCCI-2014 Tutorial

Slow Feature Analysis for Curiosity-Driven Agents

Matthew Luciw,
Varun Raj Kompella and
Juergen Schmidhuber
IDSIA, Galleria 2, Manno-Lugano 6928, Switzerland

December 16, 2013

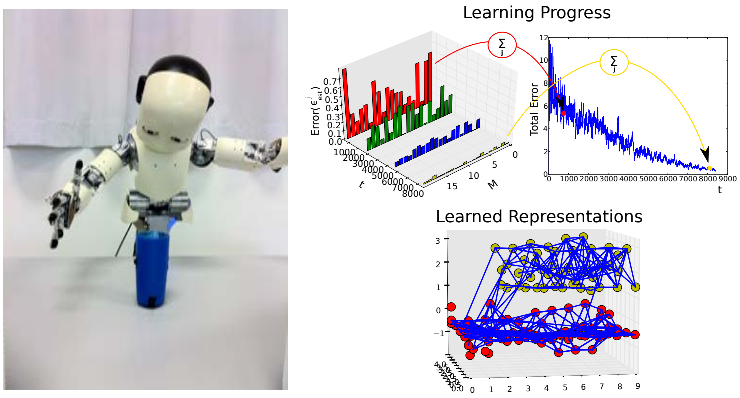

1 Goal of the Tutorial

How could a baby humanoid robot, confronted with raw-pixel high-dimensional visual input streams, learn meaningful environment representations through autonomous, real-time interactions, in the absence of a teacher? How can it leverage these learned representations to build a repertoire of skills? We present recently developed techniques and systems to tackle this problem, which share the common feature of exploitation of the *slowness* principle, via Slow Feature Analysis (SFA; [24]).

We show how incremental [3–5, 9, 10] and modular versions of SFA [2, 11] can be combined with *curiosity*-driven agents [15, 17] to enable autonomous skill learning. Curiosity-driven agents use reinforcement learning (RL; [19]) to quickly adapt control policies to maximize *intrinsic reward* (the measurable improvement in the compressor or predictor or world-model). In our case, the learning progress of the slow features is measured and used as this intrinsic reward. An agent can combine curiosity and SFA to autonomously learn behaviors, such as toppling or grasping an object, from high-dimensional image data in an unsupervised and open-ended manner [11, 20].

We use illustrative examples to present the material in detail. We present experimental results conducted on our iCub humanoid robot[1]. These results were positively received in the recently concluded IM-CLeVeR project. The intended audience is those who are interested in the intersection of Robotics and Machine Learning, and those who are interested in Developmental Robotics [12, 16, 21].

2 History

This material has not been presented in tutorial form before. Kompella presented relevant material at IJCAI (2011; [5]), Humanoids (2011; [4]) and at IM-CLeVeR project and review meetings. Luciw presented relevant material at the workshop on deep hierarchies in vision (2012; [9]), ICDL-EPIROB (2011; [8],2012; [2]), and ICANN (2012; [10]). Schmidhuber has presented artificial curiosity to large audiences in many different venues (http://www.idsia.ch/ juergen/videos.html).

3 Tutorial Description

This will be a two hour tutorial.

1. Autonomous Mental Development and Developmental Robotics
2. Artificial Curiosity
3. Slow-Feature Analysis for Raw-Pixel Input Streams
4. Incremental Slow-Feature Analysis (IncSFA)
5. AutoEncoders and IncSFA
6. IncSFA and Reinforcement Learning for Curiosity-Driven Learning
7. Curiosity-driven Modular Incremental Slow-Feature Analysis
8. Results on the iCub

An autonomous agent (i.e., a baby robot) exploring its environment encounters massive amounts of sensory data. Much of the sensorimotor data presents compressible representations in a lower-dimensional space. This representative mapping is called an abstraction [6]. A good set of abstractions makes it computationally much easier for the agent to handle data and carry out meaningful tasks, but it is an open problem how to learn such a good set, through autonomous exploration. We illustrate how such abstractions can be learned through exploration with a combination of modular incremental SFA (for the abstractions) and artificial curiosity (for the exploration).

Since the data sensed by the agent is a result of its own actions or of some time-varying external-factor, the data is often temporally correlated and therefore can be greatly compressed. Early in the tutorial, we present Slow Feature Analysis (SFA) [24]. SFA has shown much success in many problems and scenarios for extracting temporal regularities and invariant representations of the raw sensory input [1, 7, 18, 23]. However, SFA is a batch algorithm, which makes

[1] `http://www.youtube.com/watch?v=OTqdXbTEZpE`

it infeasible for open-ended learning. We present the incremental implementation of the SFA algorithm [3, 5, 9] (using Candid Covariance Free Incremental Principal Component Analysis [22] and Minor Component Analysis [13,14]) and its non-linear variant using AutoEncoders [4].

Since our focus is on learning from the environment in the absence of any external goals, the agent needs to be self-motivated, or curious. The *Formal Theory of Fun and Creativity* [15, 17] mathematically formalizes driving forces behind all kinds of curious and creative behavior. We discuss Artificial Curiosity in general. We discuss the combination of slow-feature learning and RL, which uses the learned features as basis funciotns [10]. We present the algorithm Curious Dr. MISFA [2, 11], which enables an agent, such as the iCub robot, to carry out real-time intrinsically-motivated interactions with the environment, to develop slow-features from the raw pixel inputs. CD-MISFA-driven agent undergoes coupled perceptual *and* skill learning. This is demonstrated with experimental results (with videos) conducted on the iCub humanoid robot, which uses Curious Dr. MISFA algorithm to learn skills such as toppling an object, grasping an object, and picking up the object and placing it. These skills are learned starting from low-level pixel inputs and joint angle-change outputs.

4 Prerequisite Knowledge

Participants should have some knowledge of Principal Component Analysis, Reinforcement Learning, and Autoencoder Neural Networks.

5 Biographies

5.1 Matthew Luciw

Galleria 2, Manno-Lugano 6928, Swizerland
+41 078 839 3460
matthew@idsia.ch

5.1.1 Short Bio

Matthew Luciw received the M.S. and Ph.D. degrees in Computer Science from Michigan State University (MSU) in 2006 and 2010, respectively. He completed the Interdepartmental Graduate Specialization in Cognitive Science at MSU, 2010. He worked at Microsoft Research during the summer of 2008. He is currently a postdoctoral researcher at the Swiss AI lab, IDSIA. His research interests include Autonomous Mental Development, Unsupervised Learning, Reinforcement Learning, Neural Networks, and Artificial Curiosity. He has authored or co-authored over 25 peer-reviewed research papers in well-known conferences and journals. His web page: www.idsia.ch/∼luciw

5.1.2 Teaching Experience

Teaching Assistantships at Michigan State University:

- CSE 498: Collaborative Design (Fall 2006, 2007, 2009, Spring 2007, 2008, 2009) Directed weekly project planning meetings of teams of Computer Science students working on semester-long projects with industry partners.
- CSE 260: Discrete Structures in Computer Science (Spring 2006) - Weekly lectures.
- CSE 101: Computing Concepts and Competencies (Summer 2006) - Three lectures per week.
- CSE 410: Operating Systems (Spring 2005)
- CSE 440: Introduction to Artificial Intelligence (Fall 2005)

5.2 Varun Raj Kompella

Galleria 2, Manno-Lugano 6928, Switzerland
+41-58-666 6717
varun@idsia.ch

5.2.1 Short Bio

Varun Raj Kompella is currently a researcher at the Swiss AI lab, IDSIA and a PhD candidate in the Informatics department at the Unversita della Svizzera Italiana in Lugano, Switzerland. He received his MS degree in informatics with a specialization in graphics, vision and robotics from Institut Nationale Polytechnique de Grenoble (INPG), France and carried out research at the PERCEPTION, INRIA - Rhone-Alpes, France. His work is in the area of Artificial Intelligence and he published articles in peer-reviewed conferences (IJCAI, Humanoids, ICDL) and Journals (Neural Computation, Frontiers in Neurorobotics). His web page: www.idsia.ch/∼kompella.

References

[1] M. Franzius, H. Sprekeler, and L. Wiskott. Slowness and sparseness lead to place, head-direction, and spatial-view cells. *PLoS Computational Biology*, 3(8):e166, 2007.

[2] V. R. Kompella, M. Luciw, M. Stollenga, L. Pape, and J. Schmidhuber. Autonomous learning of abstractions using curiosity-driven modular incremental slow feature analysis. In *Proc. of the Second Joint Conference on Development and Learning and Epigenetic Robotics (ICDL-EPIROB)*, San Diego, 2012. IEEE.

[3] V. R. Kompella, M. D. Luciw, and J. Schmidhuber. Incremental slow feature analysis: Adaptive low-complexity slow feature updating from high-dimensional input streams. *Neural Computation*, 24(11):2994–3024, 2012.

[4] V. R. Kompella, L. Pape, J. Masci, M. Frank, and J. Schmidhuber. Autoincsfa and vision-based developmental learning for humanoid robots. In *IEEE-RAS International Conference on Humanoid Robots*, Bled, Slovenia, 2011.

[5] V.R. Kompella, M. Luciw, and J. Schmidhuber. Incremental slow feature analysis. In *Proc. 20th International Joint Conference of Artificial Intelligence (IJCAI)*, Barcelona, 2011.

[6] G. Konidaris and A. Barto. Efficient skill learning using abstraction selection. In *Proceedings of the Twenty First International Joint Conference on Artificial Intelligence*, pages 1107–1112, 2009.

[7] R. Legenstein, N. Wilbert, and L. Wiskott. Reinforcement learning on slow features of high-dimensional input streams. *PLoS Computational Biology*, 6(8), 2010.

[8] M. Luciw, V. Graziano, M. Ring, and J. Schmidhuber. Artificial curiosity with planning for autonomous perceptual and cognitive development. In *Proc. First Joint Conference on Development Learning and on Epigenetic Robotics ICDL-EPIROB*, Frankfurt, 2011. IEEE.

[9] M. Luciw, V. R. Kompella, and J. Schmidhuber. Hierarchical incremental slow feature analysis. In *Workshop on Deep Hierarchies in Vision*, Vienna, 2012.

[10] M. Luciw and J. Schmidhuber. Low complexity proto-value function learning from sensory observations with incremental slow feature analysis. In *Proc. 22nd International Conference on Artificial Neural Networks (ICANN)*, Lausanne, 2012.

[11] Matthew Luciw, Varun Kompella, Sohrob Kazerounian, and Juergen Schmidhuber. An intrinsic value system for developing multiple invariant representations with incremental slowness learning. *Frontiers in neurorobotics*, 7, 2013.

[12] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2003.

[13] E. Oja. Principal components, minor components, and linear neural networks. *Neural Networks*, 5(6):927–935, 1992.

[14] D. Peng, Z. Yi, and W. Luo. Convergence analysis of a simple minor component analysis algorithm. *Neural Networks*, 20(7):842–850, 2007.

[15] J. Schmidhuber. Curious model-building control systems. In *Proceedings of the International Joint Conference on Neural Networks, Singapore*, volume 2, pages 1458–1463. IEEE press, 1991.

[16] J. Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187, 2006.

[17] J. Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990-2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230 –247, 2010.

[18] H. Sprekeler, T. Zito, and L. Wiskott. An extension of slow feature analysis for nonlinear blind source separation. 2010.

[19] R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press, 1998.

[20] M. Luciw V. R. Kompella, M. Stollenga and J. Schmidhuber. Continual skill learning on the icub humanoid with curious dr. misfa. *Artificial Intelligence*, 2013. Under Review.

[21] J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur, and E. Thelen. Autonomous mental development by robots and animals. *Science*, 291(5504):599–600, 2001.

[22] J. Weng, Y. Zhang, and W. Hwang. Candid covariance-free incremental principal component analysis. *Pattern Analysis and Machine Intelligence*, 25(8):1034–1040, 2003.

[23] L. Wiskott. Estimating driving forces of nonstationary time series with slow feature analysis. *Arxiv preprint cond-mat/0312317*, 2003.

[24] Laurenz Wiskott and Terrence Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.