Machine Learning for Clinicians: Advances for Multimodal Health Data

A Tutorial At MLHC 2018

Michael C. Hughes
Assistant Professor of Computer Science, Tufts University

August 17, 2018

Abstract

This is the accompanying lightly-annotated bibliography to a tutorial at the Machine Learning for Healthcare (MLHC) 2018 conference. Please see the tutorial outline and slides (which this bibliography follows): https://www.michaelchughes.com/mlhc2018_tutorial.html.

Contents

1 Overview

Tutorials targeted at clinicians:

- "Machine Learning in Medicine" (Deo, 2015) (introduces basic concepts like "supervised" and "unsupervised" learning)
- "Introduction to Machine Learning" (written for audience of Methods in Molecular Biology journal) (Baştanlar & Özuysal, 2014)

Accessible ML Textbooks for Practioners

- "Evaluating Machine Learning Models" (Zheng, 2015)

Calls to Action:

- "Opportunities for Machine Learning in Healthcare": (Ghassemi et al., 2018)
- "Machine Learning that Matters" (Wagstaff, 2012)
- "What this Computer Needs is a Physician" (Verghese et al., 2018)

Highlighted recent methods:

- MGP-RNN for Sepsis Risk Prediction (Futoma et al., 2017)

Surveys:

- Survey: Deep Learning for EHR in JAMIA by Xiao et al.
- Survey: "Opportunities and obstacles for deep learning in biology and medicine" by Ching et al. (2018)
- Survey: Deep Learning for Medical Imaging by Litjens et al. (2017)

[Baştanlar & Özuysal 2014] BAŞTANLAR, Yalin ; ÖZUYSAL, Mustafa: Introduction to Machine Learning. In: *miRNomics: MicroRNA Biology and Computational Analysis*. Humana Press, Totowa, NJ, 2014 (Methods in Molecular Biology), p. 105–128. – ISBN 978-1-62703-747-1 978-1-62703-748-8

[Ching et al. 2018] CHING, Travers ; HIMMELSTEIN, Daniel S. ; BEAULIEU-JONES, Brett K. ; KALININ, Alexandr A. ; DO, Brian T. ; WAY, Gregory P. ; FERRERO, Enrico ; AGAPOW, Paul-Michael ; ZIETZ, Michael ; HOFFMAN, Michael M. ; XIE, Wei ; ROSEN, Gail L. ; LENGERICH, Benjamin J. ; ISRAELI, Johnny ; LANCHANTIN, Jack ; WOLOSZYNEK, Stephen ; CARPENTER, Anne E. ; SHRIKUMAR, Avanti ; XU, Jinbo ; COFER, Evan M. ; LAVENDER, Christopher A. ; TURAGA, Srinivas C. ; ALEXANDARI, Amr M. ; LU, Zhiyong ; HARRIS, David J. ; DECAPRIO, Dave ; QI, Yanjun ; KUNDAJE, Anshul ; PENG, Yifan ; WILEY, Laura K. ; SEGLER, Marwin H. S. ; BOCA, Simina M. ; SWAMIDASS, S. J. ; HUANG, Austin ; GITTER, Anthony ; GREENE, Casey S.: Opportunities and Obstacles for Deep Learning in Biology and Medicine. In: *Journal of the Royal Society, Interface* 15 (2018), Nr. 141. – ISSN 1742-5662

[Deo 2015] DEO, R. C.: Machine Learning in Medicine. In: *Circulation* 132 (2015), Nr. 20, p. 1920–1930. – ISSN 0009-7322, 1524-4539

[Futoma et al. 2017] FUTOMA, Joseph ; HARIHARAN, Sanjay ; SENDAK, Mark ; BRAJER, Nathan ; CLEMENT, Meredith ; BEDOYA, Armando ; O'BRIEN, Cara ; HELLER, Katherine: An Improved Multi-Output Gaussian Process RNN with Real-Time Validation for Early Sepsis Detection. In: *arXiv:1708.05894 [Stat]*, URL http://arxiv.org/abs/1708.05894, 2017

[Ghassemi et al. 2018] GHASSEMI, Marzyeh ; NAUMANN, Tristan ; SCHULAM, Peter ; BEAM, Andrew L. ; RANGANATH, Rajesh: Opportunities in Machine Learning for Healthcare. In: *arXiv:1806.00388 [cs, stat]* (2018)

[Litjens et al. 2017] LITJENS, Geert ; KOOI, Thijs ; BEJNORDI, Babak E. ; SETIO, Arnaud Arindra A. ; CIOMPI, Francesco ; GHAFOORIAN, Mohsen ; VAN DER LAAK, Jeroen A. W. M. ; VAN GINNEKEN, Bram ; SÁNCHEZ, Clara I.: A Survey on Deep Learning in Medical Image Analysis. In: *Medical Image Analysis* 42 (2017), p. 60–88. – ISSN 1361-8423

[Verghese et al. 2018] VERGHESE, Abraham ; SHAH, Nigam H. ; HARRINGTON, Robert A.: What This Computer Needs Is a Physician: Humanism and Artificial Intelligence. In: *JAMA* 319 (2018), Nr. 1, p. 19–20. – ISSN 1538-3598

[Wagstaff 2012] WAGSTAFF, Kiri: Machine Learning That Matters. In: *arXiv:1206.4656 [cs, stat]* (2012)

[Xiao et al.] XIAO, Cao ; CHOI, Edward ; SUN, Jimeng: Opportunities and Challenges in Developing Deep Learning Models Using Electronic Health Records Data: A Systematic Review. In: *Journal of the American Medical Informatics Association*

[Zheng 2015] ZHENG, Alice: *Evaluating Machine Learning Models: A Beginner's Guide to Key Concepts and Pitfalls*. O'Reilly, 2015

2 Making and Evaluating Predictions

Evaluating Predictions

Dividing Data into Train/Test/Validation sets and Cross-validation:

- Ch. 7 of (Hastie et al., 2009)
- (Breiman & Spector, 1992)

Evaluating binary classifiers :

- See Zheng (2015)
- "Evaluation of binary classifiers" on Wikipedia fo formulas https://en.wikipedia.org/wiki/Evaluation_of_binary_classifiers.

Intro to ROC analysis: (Fawcett, 2006).

Limitations of Area under ROC curve: (Romero-Brufau et al., 2015) and (Hand, 2009). Also see this blog post by Luke Oakden-Rayner https://lukeoakdenrayner.wordpress.com/2018/01/07/the-philosophical-argument-for-using-roc-curves/

Utility analysis using fixed costs for TP, FP, TN, FN: Blog post by Nicholas Krutchen http://blog.mldb.ai/blog/posts/2016/01/ml-meets-economics/

Setting a decision threshold: (Irwin & Irwin, 2011)

Cost curves: (Drummond & Holte, 2006)

Decision curve analysis (Rousson & Zumbrunn, 2011) and (Vickers & Elkin, 2006)

Best practices for model evaluation : (Steyerberg & Vergouwe, 2014)

Making Predictions

Linear Regression: Ch. 3 of (Hastie et al., 2009)

Logistic Regression: Ch 4.4 of (Hastie et al., 2009)

Decision trees : Ch. 2.9 of (Hastie et al., 2009)

Random forests : Ch 15 of (Hastie et al., 2009)

Hyperparameter tuning : See the (unnumbered) chapter of coverage in (Zheng, 2015)

Gaussian processes: (Rasmusen & Williams, 2006)

[Breiman & Spector 1992] BREIMAN, Leo ; SPECTOR, Philip: Submodel Selection and Evaluation in Regression. The X-Random Case. In: *International Statistical Review / Revue Internationale de Statistique* 60 (1992), Nr. 3, p. 291–319. – ISSN 0306-7734

[Drummond & Holte 2006] DRUMMOND, Chris ; HOLTE, Robert C.: Cost Curves: An Improved Method for Visualizing Classifier Performance. In: *Machine Learning* 65 (2006), Nr. 1, p. 95–130. – ISSN 0885-6125, 1573-0565

[Fawcett 2006] FAWCETT, Tom: An Introduction to ROC Analysis. In: *Pattern Recognition Letters* 27 (2006), Nr. 8, p. 861–874. – ISSN 0167-8655

[Hand 2009] HAND, David J.: Measuring Classifier Performance: A Coherent Alternative to the Area under the ROC Curve. In: *Machine Learning* 77 (2009), Nr. 1, p. 103–123. – ISSN 0885-6125, 1573-0565

[Hastie et al. 2009] HASTIE, Trevor ; TIBSHIRANI, Robert ; FRIEDMAN, Jerome: *The Elements of Statistical Learning*. Second. URL https://web.stanford.edu/~hastie/ElemStatLearn/, 2009 (Springer Series in Statistics)

[Irwin & Irwin 2011] IRWIN, R. J. ; IRWIN, Timothy C.: A Principled Approach to Setting Optimal Diagnostic Thresholds: Where ROC and Indifference Curves Meet. In: *European Journal of Internal Medicine* 22 (2011), Nr. 3, p. 230–234. – ISSN 1879-0828

[Pedregosa et al. 2011] PEDREGOSA, Fabian ; VAROQUAUX, Gaël ; GRAMFORT, Alexandre ; MICHEL, Vincent ; THIRION, Bertrand ; GRISEL, Olivier ; BLONDEL, Mathieu ; PRETTENHOFER, Peter ; WEISS, Ron ; DUBOURG, Vincent ; VANDERPLAS, Jake ; PASSOS, Alexandre ; COURNAPEAU, David ; BRUCHER, Matthieu ; PERROT, Matthieu ; DUCHESNAY, Édouard: Scikit-Learn: Machine Learning in Python. In: *Journal of Machine Learning Research* 12 (2011), p. 2825–2830. – http://scikit-learn.org/stable/

[Rasmusen & Williams 2006] RASMUSEN, Carl E. ; WILLIAMS, Christopher K I.: Gaussian Processes for Machine Learning. In: *Gaussian Processes for Machine Learning*. URL http://www.gaussianprocess.org/gpml/, 2006, p. Ch. 2

[Romero-Brufau et al. 2015] ROMERO-BRUFAU, Santiago ; HUDDLESTON, Jeanne M. ; ESCOBAR, Gabriel J. ; LIEBOW, Mark: Why the C-Statistic Is Not Informative to Evaluate Early Warning Scores and What Metrics to Use. In: *Critical Care* 19 (2015), Nr. 1. – ISSN 1364-8535

[Rousson & Zumbrunn 2011] ROUSSON, Valentin ; ZUMBRUNN, Thomas: Decision Curve Analysis Revisited: Overall Net Benefit, Relationships to ROC Curve Analysis, and Application to Case-Control Studies. In: *BMC Medical Informatics and Decision Making* 11 (2011), p. 45. – ISSN 1472-6947

[Steyerberg & Vergouwe 2014] STEYERBERG, Ewout W. ; VERGOUWE, Yvonne: Towards Better Clinical Prediction Models: Seven Steps for Development and an ABCD for Validation. In: *European Heart Journal* 35 (2014), Nr. 29, p. 1925–1931. – ISSN 0195-668X

[Vickers & Elkin 2006] VICKERS, Andrew J. ; ELKIN, Elena B.: Decision Curve Analysis: A Novel Method for Evaluating Prediction Models. In: *Medical decision making : an international journal of the Society for Medical Decision Making* 26 (2006), Nr. 6, p. 565–574. – ISSN 0272-989X

[Zheng 2015] ZHENG, Alice: *Evaluating Machine Learning Models: A Beginner's Guide to Key Concepts and Pitfalls*. O'Reilly, 2015

3 Learning Representations

Bag-of-words Representations

Topic Models

- Topic models survey (Blei, 2012)

Tensor Factorization/Topic Models for EHR

- Marble (Ho et al., 2014b)
- Limestone (Ho et al., 2014a)
- TaGiTeD (Yang et al., 2017)
- PC-sLDA (Hughes et al., 2018)

Learned Image Representations

Convolutional Neural Networks

- Deep CNNs for ImageNet (Krizhevsky et al., 2012)
- https://www.tensorflow.org/tutorials/images/deep_cnn

Learned Time Series Representations

Highlighted ML+Health Papers

- "Learning to Diagnose with LSTMs" (Lipton et al., 2015)

Hidden Markov Models that do not require aligned time series

- (Liu et al., 2015)
- (Leiva-murillo et al., 2011)

Learned Text Representations

Bidirectional LSTMs:

- (Schuster & Paliwal, 1997)
- (Graves & Schmidhuber, 2005)

1D Convolutional NNs (Zhang & Wallace, 2015)

Word Embeddings

- GloVe (Pennington et al., 2014)
- word2vec (Mikolov et al., 2013)
- med2vec for EHR codes (Choi et al., 2016)
- Applied to Radiology Report text: (Banerjee et al., 2018)

Tricks of the Trade

Dropout (Srivastava et al., 2014)

Data Augmentation Example for Melanoma Classification (Vasconcelos & Vasconcelos, 2017)

Target/Label Replication Example of LSTM adding loss signal to each timestep, not just final one: (Lipton et al., 2015)

Models that generate data

Denoising Autoencoders

- Denoising AEs (Vincent et al., 2008)
- Deep Patient (Miotto et al., 2016)

Deep generative models and Variational autoencoders: Johnson et al. (2016) and Kingma & Welling (2014)

GANs: Goodfellow et al. (2014)

medGAN: Choi et al. (2016)

[Banerjee et al. 2018] BANERJEE, Imon ; CHEN, Matthew C. ; LUNGREN, Matthew P. ; RUBIN, Daniel L.: Radiology Report Annotation Using Intelligent Word Embeddings: Applied to Multi-Institutional Chest CT Cohort. In: *Journal of Biomedical Informatics* 77 (2018), p. 11–20. – ISSN 1532-0464

[Blei 2012] BLEI, David M.: Probabilistic Topic Models. In: *Communications of the ACM* 55 (2012), Nr. 4, p. 77–84

[Choi et al. 2016] CHOI, Edward ; BAHADORI, Mohammad T. ; SEARLES, Elizabeth ; COFFEY, Catherine ; THOMPSON, Michael ; BOST, James ; TEJEDOR-SOJO, Javier ; SUN, Jimeng: Multi-Layer Representation Learning for Medical Concepts. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*. San Francisco, California, USA : ACM Press, 2016, p. 1495–1504. – ISBN 978-1-4503-4232-2

[Goodfellow et al. 2014] GOODFELLOW, Ian J. ; POUGET-ABADIE, Jean ; MIRZA, Mehdi ; XU, Bing ; WARDE-FARLEY, David ; OZAIR, Sherjil ; COURVILLE, Aaron ; BENGIO, Yoshua: Generative Adversarial Networks. In: *Advances in Neural Information Processing Systems*, URL http://arxiv.org/abs/1406.2661, 2014

[Graves & Schmidhuber 2005] GRAVES, Alex ; SCHMIDHUBER, Jürgen: Framewise Phoneme Classification with Bidirectional LSTM and Other Neural Network Architectures. In: *Neural Networks* 18 (2005), Nr. 5, p. 602–610. – ISSN 0893-6080

[Ho et al. 2014a] HO, Joyce C. ; GHOSH, Joydeep ; STEINHUBL, Steve R. ; STEWART, Walter F. ; DENNY, Joshua C. ; MALIN, Bradley A. ; SUN, Jimeng: Limestone: High-Throughput Candidate Phenotype Generation via Tensor Factorization. In: *Journal of Biomedical Informatics* 52 (2014), p. 199–211. – ISSN 15320464

[Ho et al. 2014b] HO, Joyce C. ; GHOSH, Joydeep ; SUN, Jimeng: Marble: High-Throughput Phenotyping from Electronic Health Records via

Sparse Nonnegative Tensor Factorization, ACM Press, 2014, p. 115–124. – ISBN 978-1-4503-2956-9

[Hughes et al. 2018] HUGHES, Michael C. ; HOPE, Gabriel ; WEINER, Leah ; MCCOY, Thomas H. ; PERLIS, Roy H. ; SUDDERTH, Erik B. ; DOSHI-VELEZ, Finale: Semi-Supervised Prediction-Constrained Topic Models. In: *Artificial Intelligence and Statistics*, 2018

[Johnson et al. 2016] JOHNSON, Matthew ; DUVENAUD, David K. ; WILTSCHKO, Alex ; ADAMS, Ryan P. ; DATTA, Sandeep R.: Composing Graphical Models with Neural Networks for Structured Representations and Fast Inference. In: LEE, D. D. (Editor) ; SUGIYAMA, M. (Editor) ; LUXBURG, U. V. (Editor) ; GUYON, I. (Editor) ; GARNETT, R. (Editor): *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2016, p. 2946–2954

[Kingma & Welling 2014] KINGMA, Diederik P. ; WELLING, Max: Auto-Encoding Variational Bayes. In: *International Conference on Learning Representations*, URL http://arxiv.org/abs/1312.6114, 2014

[Krizhevsky et al. 2012] KRIZHEVSKY, Alex ; SUTSKEVER, Ilya ; HINTON, Geoffrey E.: ImageNet Classification with Deep Convolutional Neural Networks. In: *Advances in Neural Information Processing Systems*, URL https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf, 2012

[Leiva-murillo et al. 2011] LEIVA-MURILLO, Jose M. ; RODRÍGUEZ, Antonio A. ; BACA-GARCÍA, Enrique ; DÍAZ, Fundación J.: Visualization and Prediction of Disease Interactions with Continuous-Time Hidden Markov Models, URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.422.8632&rep=rep1&type=pdf, 2011

[Lipton et al. 2015] LIPTON, Zachary C. ; KALE, David C. ; ELKAN, Charles ; WETZEL, Randall: Learning to Diagnose with LSTM Recurrent Neural Networks. In: *arXiv:1511.03677 [Cs]*, URL http://arxiv.org/abs/1511.03677, 2015

[Liu et al. 2015] LIU, Yu-Ying ; LI, Shuang ; LI, Fuxin ; SONG, Le ; REHG, James M.: Efficient Learning of Continuous-Time Hidden Markov Models for Disease Progression. In: *Advances in neural information processing systems* 28 (2015), p. 3599–3607. – ISSN 1049-5258

[Mikolov et al. 2013] MIKOLOV, Tomas ; CHEN, Kai ; CORRADO, Greg ; DEAN, Jeffrey: Efficient Estimation of Word Representations in Vector Space. In: *arXiv:1301.3781 [cs]* (2013)

[Miotto et al. 2016] MIOTTO, Riccardo ; LI, Li ; KIDD, Brian A. ; DUDLEY, Joel T.: Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. In: *Scientific Reports* 6 (2016), p. 26094. – ISSN 2045-2322

[Pennington et al. 2014] PENNINGTON, Jeffrey ; SOCHER, Richard ; MANNING, Christopher: Glove: Global Vectors for Word Representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar : Association for Computational Linguistics, 2014, p. 1532–1543

[Schuster & Paliwal 1997] SCHUSTER, M. ; PALIWAL, K.K.: Bidirectional Recurrent Neural Networks. In: *Trans. Sig. Proc.* 45 (1997), Nr. 11, p. 2673–2681. – ISSN 1053-587X

[Srivastava et al. 2014] SRIVASTAVA, Nitish ; HINTON, Geoffrey ; KRIZHEVSKY, Alex ; SUTSKEVER, Ilya ; SALAKHUTDINOV, Ruslan: Dropout: A Simple Way to Prevent Neural Networks from Overfitting. In: *Journal of Machine Learning Research* 15 (2014), p. 30

[Vasconcelos & Vasconcelos 2017] VASCONCELOS, Cristina N. ; VASCONCELOS, Bárbara N.: Convolutional Neural Network Committees for Melanoma Classification with Classical And Expert Knowledge Based Image Transforms Data Augmentation. In: *arXiv:1702.07025 [cs]* (2017)

[Vincent et al. 2008] VINCENT, Pascal ; LAROCHELLE, Hugo ; BENGIO, Yoshua ; MANZAGOL, Pierre-Antoine: Extracting and Composing Robust Features with Denoising Autoencoders, ACM Press, 2008, p. 1096–1103. – ISBN 978-1-60558-205-4

[Yang et al. 2017] YANG, Kai ; LI, Xiang ; LIU, Haifeng ; MEI, Jing ; XIE, Guotong ; ZHAO, Junfeng ; XIE, Bing ; WANG, Fei: TaGiTeD: Predictive

Task Guided Tensor Decomposition for Representation Learning from Electronic Health Records. In: *AAAI Conference on Artificial Intelligence*, 2017, p. 7

[Zhang & Wallace 2015] ZHANG, Ye ; WALLACE, Byron: A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification. In: *arXiv:1510.03820 [cs]* (2015)

4 Missing Data

Motivating example:

- Time-of-day of lab tests and 3-year survival rate: (Agniel et al., 2018)

Highlighted Methods:

- MissForest: Random Forest for Imputing Missing data (Stekhoven & Bühlmann, 2012)
- GRU-D: RNNs that handle missingness (Che et al., 2018)
- GAIN: Generative Adversarial Imputation Networks (Yoon et al., 2018)

Other methods

- Generative model that "integrates away" missing data (Caballero Barajas & Akella, 2015)
- (Tresp & Briegel, 1997)

[Agniel et al. 2018] AGNIEL, Denis ; KOHANE, Isaac S. ; WEBER, Griffin M.: Biases in Electronic Health Record Data Due to Processes within the Healthcare System: Retrospective Observational Study. In: *BMJ* 361 (2018), Nr. k1479

[Caballero Barajas & Akella 2015] CABALLERO BARAJAS, Karla L. ; AKELLA, Ram: Dynamically Modeling Patient's Health State from Electronic Medical Records: A Time Series Approach, ACM Press, 2015, p. 69–78. – ISBN 978-1-4503-3664-2

[Che et al. 2018] CHE, Zhengping ; PURUSHOTHAM, Sanjay ; CHO, Kyunghyun ; SONTAG, David ; LIU, Yan: Recurrent Neural Networks for Multivariate Time Series with Missing Values. In: *Scientific Reports* 8 (2018), Nr. 1. – ISSN 2045-2322

[Stekhoven & Bühlmann 2012] STEKHOVEN, Daniel J. ; BÜHLMANN, Peter: MissForest—Non-Parametric Missing Value Imputation for Mixed-Type Data. In: *Bioinformatics* 28 (2012), Nr. 1, p. 112–118. – ISSN 1367-4803

[Tresp & Briegel 1997] TRESP, Volker ; BRIEGEL, Thomas: A Solution for Missing Data in Recurrent Neural Networks with an Application to Blood Glucose Prediction. In: *Advances in Neural Information Processing Systems*, 1997

[Yoon et al. 2018] YOON, Jinsung ; JORDON, James ; VAN DER SCHAAR, Mihaela: GAIN: Missing Data Imputation Using Generative Adversarial Nets. In: *International Conference on Machine Learning*, URL http://proceedings.mlr.press/v80/yoon18a.html, 2018

5 Semi-supervised Prediction

Methods that can combine few labeled examples with many unlabeled examples.

Evaluation best practices paper

- Realistic Evaluation of SSL (for images) (Oliver et al., 2018)

Highlighted specific methods

- Denoising Autoencoders for 2-stage SSL in EHR (Beaulieu-Jones & Greene, 2016)

Other interesting methods

- Semisupervised with GANs: (McDermott et al., 2018)
- Prediction-constrained training. Longer arXiv version (Hughes et al., 2017)
- Cotraining (Blum & Mitchell, 1998)
- Bayesian co-training and active sensing (given patient demographics data, which one should I image to learn the most): (Yu et al., 2011)

[Beaulieu-Jones & Greene 2016] BEAULIEU-JONES, Brett K. ; GREENE, Casey S.: Semi-Supervised Learning of the Electronic Health Record for Phenotype Stratification. In: *Journal of Biomedical Informatics* 64 (2016), p. 168–178. – ISSN 15320464

[Blum & Mitchell 1998] BLUM, Avrim ; MITCHELL, Tom: Combining Labeled and Unlabeled Data with Co-Training. In: *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*. New York, NY, USA : ACM, 1998 (COLT' 98), p. 92–100. – ISBN 978-1-58113-057-7

[Hughes et al. 2017] HUGHES, Michael C. ; WEINER, Leah ; HOPE, Gabriel ; MCCOY, Thomas H. ; PERLIS, Roy H. ; SUDDERTH, Erik B. ; DOSHI-VELEZ, Finale: Prediction-Constrained Training for Semi-Supervised Mixture and Topic Models. In: *arXiv preprint 1707.07341* (2017)

[McDermott et al. 2018] MCDERMOTT, Matthew B. A. ; YAN, Tom ; NAUMANN, Tristan ; HUNT, Nathan ; SURESH, Harini ; SZOLOVITS, Peter ; GHASSEMI, Marzyeh: Semi-Supervised Biomedical Translation With Cycle Wasserstein Regression GANs. In: *Thirty-Second AAAI Conference on Artificial Intelligence*, URL https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16938, 2018

[Oliver et al. 2018] OLIVER, Avital ; ODENA, Augustus ; RAFFEL, Colin ; CUBUK, Ekin D. ; GOODFELLOW, Ian J.: Realistic Evaluation of Deep Semi-Supervised Learning Algorithms. In: *arXiv:1804.09170 [cs, stat]* (2018)

[Yu et al. 2011] YU, Shipeng ; KRISHNAPURAM, Balaji ; ROSALES, R\'{o}mer ; RAO, R. B.: Bayesian Co-Training. In: *Journal of Machine Learning Research* 12 (2011), p. 2649–2680

6 Multimodal Prediction

How can we combine images, text, and other modalities of information to develop good learned representations?

Overviews/Surveys on ML methods

- Baltrusaitis, Ahuja, and Morency's survey: "Multimodal Machine Learning: A Survey and Taxonomy" (focus on images, text, and some video) (Baltrušaitis et al., 2017)
- Slidedeck from ACL 2017 tutorial by Morency and Baltrusaitis (accompanies paper above): https://www.cs.cmu.edu/~morency/MMML-Tutorial-ACL2017.pdf
- Another survey: (Ramachandram & Taylor, 2017)

Highlighted ML Methods Papers

- Coordinated embeddings of images and text (vector math with pictures and text) (Kiros et al., 2014)

ML+Health Examples

- Deep Poisson Factor Analysis for Multiple Types of EHR codes (medications, procedures, diagnoses) (Henao et al., 2015)
- MR and PET images for Alzheimer's: (Lu et al., 2018)
- Cervical cancer images + demographics:(Xu et al., 2016)

[Baltrušaitis et al. 2017] BALTRUŠAITIS, Tadas ; AHUJA, Chaitanya ; MORENCY, Louis-Philippe: Multimodal Machine Learning: A Survey and Taxonomy. In: *arXiv:1705.09406 [cs]* (2017)

[Henao et al. 2015] HENAO, Ricardo ; LU, James T. ; LUCAS, Joseph E. ; FERRANTI, Jeffrey ; CARIN, Lawrence: Electronic Health Record Analysis via Deep Poisson Factor Models. In: *Journal of Machine Learning Research* (2015), p. 31

[Kiros et al. 2014] KIROS, Ryan ; SALAKHUTDINOV, Ruslan ; ZEMEL, Richard S.: Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models. In: *Bayesian Deep Learning Workshop at NIPS*, URL http://arxiv.org/abs/1411.2539, 2014

[Lu et al. 2018] LU, Donghuan ; POPURI, Karteek ; DING, Gavin W. ; BALACHANDAR, Rakesh ; BEG, Mirza F.: Multimodal and Multiscale Deep Neural Networks for the Early Diagnosis of Alzheimer's Disease Using Structural MR and FDG-PET Images. In: *Scientific Reports* 8 (2018), Nr. 1, p. 5697. – ISSN 2045-2322

[Ramachandram & Taylor 2017] RAMACHANDRAM, D. ; TAYLOR, G. W.: Deep Multimodal Learning: A Survey on Recent Advances and Trends. In: *IEEE Signal Processing Magazine* 34 (2017), Nr. 6, p. 96–108. – ISSN 1053-5888

[Xu et al. 2016] XU, Tao ; ZHANG, Han ; HUANG, Xiaolei ; ZHANG, Shaoting ; METAXAS, Dimitris N.: Multimodal Deep Learning for Cervical Dysplasia Diagnosis. In: OURSELIN, Sebastien (Editor) ; JOSKOWICZ, Leo (Editor) ; SABUNCU, Mert R. (Editor) ; UNAL, Gozde (Editor) ; WELLS, William (Editor): *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016* Volume 9901. Cham : Springer International Publishing, 2016, p. 115–123. – ISBN 978-3-319-46722-1 978-3-319-46723-8

7 Interpretable Prediction

Position papers:

- Doshi-Velez & Kim (2017)
- Lipton (2016)

Classic papers on ML Interpretability in Health:

- (Caruana et al., 2015)

Highlighted Methods:

- SLIM (Ustun & Rudin, 2016)
- LIME (Ribeiro et al., 2016)
- Tree Regularization (Wu et al., 2018)

[Caruana et al. 2015] CARUANA, Rich ; LOU, Yin ; GEHRKE, Johannes ; KOCH, Paul ; STURM, Marc ; ELHADAD, Noemie: Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-Day Readmission. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*. Sydney, NSW, Australia : ACM Press, 2015, p. 1721–1730. – ISBN 978-1-4503-3664-2

[Doshi-Velez & Kim 2017] DOSHI-VELEZ, Finale ; KIM, Been: Towards A Rigorous Science of Interpretable Machine Learning. In: *arXiv:1702.08608 [cs, stat]* (2017)

[Lipton 2016] LIPTON, Zachary C.: The Mythos of Model Interpretability. In: *arXiv:1606.03490 [Cs, Stat]*, URL http://arxiv.org/abs/1606.03490, 2016

[Ribeiro et al. 2016] RIBEIRO, Marco T. ; SINGH, Sameer ; GUESTRIN, Carlos: "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, URL http://www.kdd.org/kdd2016/papers/files/rfp0573-ribeiroA.pdf, 2016

[Ustun & Rudin 2016] USTUN, Berk ; RUDIN, Cynthia: Supersparse Linear Integer Models for Optimized Medical Scoring Systems. In: *Machine Learning* 102 (2016), Nr. 3, p. 349–391. – ISSN 0885-6125, 1573-0565

[Wu et al. 2018] WU, Mike ; HUGHES, Michael C. ; PARBHOO, Sonali ; ZAZZI, Maurizio ; ROTH, Volker ; DOSHI-VELEZ, Finale: Beyond Sparsity: Tree Regularization of Deep Models for Interpretability. In: *AAAI Conference on Artificial Intelligence*, URL https://arxiv.org/pdf/1711.06178.pdf, 2018

8 Causality

Book The Book of Why, by Judea Pearl and Dana Mackenzie http://bayes.cs.ucla.edu/WHY/

Position Paper

- Pearl on why Supervised Learning isn't (and won't be) enough for causal reasoning (Pearl, 2018)

Tutorials

- "Causal Inference for Observational Studies" at ICML 2017 https://cs.nyu.edu/~shalit/tutorial.html

Highlighted papers:

- Counterfactual GP: (Schulam & Saria, 2017)
- Causal-Effect Variational Autoencoder: (Louizos et al., 2017)

[Louizos et al. 2017] LOUIZOS, Christos ; SHALIT, Uri ; MOOIJ, Joris ; SONTAG, David ; ZEMEL, Richard ; WELLING, Max: Causal Effect Inference with Deep Latent-Variable Models. In: *arXiv:1705.08821 [cs, stat]* (2017)

[Pearl 2018] PEARL, Judea: Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution. In: *arXiv:1801.04016 [cs, stat]* (2018)

[Schulam & Saria 2017] SCHULAM, Peter ; SARIA, Suchi: Reliable Decision Support Using Counterfactual Models. In: GUYON, I. (Editor) ; LUXBURG, U. V. (Editor) ; BENGIO, S. (Editor) ; WALLACH, H. (Editor) ; FERGUS, R. (Editor) ; VISHWANATHAN, S. (Editor) ; GARNETT, R. (Editor): *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2017, p. 1697–1708

9 Reinforcement Learning

Emerging best practices for RL in healthcare are covered in (Gottesman et al., 2018)

Highlighted papers applying RL to real sequential treatment problems in healthcare:

- RL for Sepsis Treatment: (Raghu et al., 2017)
- RL for Schizophrenia: (Shortreed et al., 2011)
- RL for Mechanical Ventilation: (Prasad et al., 2017)

[Gottesman et al. 2018] GOTTESMAN, Omer ; JOHANSSON, Fredrik ; MEIER, Joshua ; DENT, Jack ; LEE, Donghun ; SRINIVASAN, Srivatsan ; ZHANG, Linying ; DING, Yi ; WIHL, David ; PENG, Xuefeng ; YAO, Jiayu ; LAGE, Isaac ; MOSCH, Christopher ; LEHMAN, Li-wei H. ; KOMOROWSKI, Matthieu ; KOMOROWSKI, Matthieu ; FAISAL, Aldo ; CELI, Leo A. ; SONTAG, David ; DOSHI-VELEZ, Finale: Evaluating Reinforcement Learning Algorithms in Observational Health Settings. In: *arXiv:1805.12298 [cs, stat]* (2018)

[Prasad et al. 2017] PRASAD, Niranjani ; CHENG, Li-Fang ; CHIVERS, Corey ; DRAUGELIS, Michael ; ENGELHARDT, Barbara E.: A Reinforcement Learning Approach to Weaning of Mechanical Ventilation in Intensive Care Units. In: *Uncertainty in Artificial Intelligence*, URL http://auai.org/uai2017/proceedings/papers/209.pdf, 2017, p. 10

[Raghu et al. 2017] RAGHU, Aniruddh ; KOMOROWSKI, Matthieu ; CELI, Leo A. ; SZOLOVITS, Peter ; GHASSEMI, Marzyeh: Continuous State-Space Models for Optimal Sepsis Treatment: A Deep Reinforcement Learning Approach. In: *Machine Learning for Healthcare Conference*, URL http://proceedings.mlr.press/v68/raghu17a.html, 2017, p. 147–163

[Shortreed et al. 2011] SHORTREED, Susan M. ; LABER, Eric ; LIZOTTE, Daniel J. ; STROUP, T. S. ; PINEAU, Joelle ; MURPHY, Susan A.: Informing Sequential Clinical Decision-Making through Reinforcement Learning: An Empirical Study. In: *Machine learning* 84 (2011), Nr. 1-2, p. 109–136. – ISSN 0885-6125