HUMAN-LEVEL CONCEPT LEARNING THROUGH PROBABILISTIC PROGRAM INDUCTION pdf

1: Human-level concept learning through probabilistic program induction : MachineLearning

RESEARCH ARTICLES COGNITIVE SCIENCE Human-level concept learning through probabilistic program induction Brenden M. Lake,1 Ruslan Salakhutdinov,2 Joshua B. Tenenbaum3 People learning new concepts can often generalize successfully from just a single example.*

This paper was formerly titled "Elemental causal induction. Discovering latent classes in relational data. Semi-supervised learning with trees. Advances in Neural Information Processing Systems Hierarchical topic models and the nested Chinese restaurant process. Backwards blocking and Bayesian reasoning in preschoolers. Gopnik , Cognitive Science, 28, Using physical theories to infer hidden causal structure. The role of causal models in reasoning under uncertainty. V1 neurons signal acquisition of an internal representation of stimulus location. Probability, algorithmic complexity, and subjective randomness. Cambridge, MIT Press, , Inferring causal networks from observations and interventions. Blum , Cognitive Science Science Jan 4 Global versus local methods in nonlinear dimensionality reduction. Unsupervised learning of curved manifolds. Bayesian models of inductive generalization. Griffiths , Behavioral and Brain Sciences, 24 pp. Some specifics about generalization. Griffiths , Behavioral and Brain Sciences, 24, pages The rational basis of representativeness. Reconciling intuition and probability theory. Article in Psychology Today single-page view Structure learning in human causal induction. De Silva and J. Science , Website Separating style and content with bilinear models. Neural Computation 12 6 , Xu , Proceedings of the 22nd Annual Conference of the Cognitive Science Society postscript Teacakes, trains, toxins, and taxicabs: A Bayesian account of predicting the future. Thesis, MIT, Mapping a manifold of perceptual observations. The Isomap Algorithm and Topological Stability. A global geometric framework for nonlinear dimensionality reduction. Website Mapping a manifold of perceptual observations. Integrating topics and syntax. The large-scale structure of semantic networks: Tenenbaum , Cognitive Science, 29 1. Xu , Proceedings of the 22nd Annual Conference of the Cognitive Science Society postscript Learning and similarity A generative theory of similarity. Generalization, similarity, and Bayesian inference. Learning the structure of similarity. Article in The Economist From algorithmic to subjective randomness. Teacakes, trains, toxins, and taxicabs: Psychology, philosophy, and computation. Two proposals for causal grammars.

HUMAN-LEVEL CONCEPT LEARNING THROUGH PROBABILISTIC PROGRAM INDUCTION pdf

2: Human-level concept learning through probabilistic program induction. - Abstract - Europe PMC

Learning proceeds by constructing programs that best explain the observations under a Bayesian criterion, and the model "learns to learn" (23, 24) by developing hierarchical priors that allow previous experience with related concepts to ease learning of new concepts (25, 26).

Viewed from a computer vision perspective, HBPL is a drastic departure from the deep convolutional neural networks that dominate the field, and while HBPLs may not be ready for production systems, the remarkable accuracy and careful composition of concept learning provide insights into how humans and machines learn. The two papers NIPS and Science versions build on a long tradition of probabilistic programming efforts, and also on advances in parsing and modelling hand-written symbols. The HBPL model can also be viewed as an extension of other structural description models. The two papers discussed here make a threefold contribution. Second, they propose the HBPL algorithm, which is compositional, generative and causal. The HBPL model itself is rather sophisticated, with a plethora of parameters. Indeed, the authors could have made good use of the richness of Omniglot in denoting the various parameters, relations, and hyper-parameters. At the core of HBPL lies the concept of character types e. A character type is modelled by a 3-tuple of stroke-count, strokes, and relations. Strokes are in-turn broken down into sub-strokes, with intuitive sub-stroke splits occurring when the pen rests or drastically changes direction. Sub-strokes are in-turn defined by a tuple of: The third component of the stroke, the Relation, define relations between sub-strokes, e. The Tokens are conceptually broken down into pen trajectories and an ink model, each with its own generative model and hyper-parameters. Learning is done in two stages. The remaining 20 alphabets are used for one-shot learning and sample synthesis experiments. On a high level, inference uses bayes classification rule to find the character type that best explains the given input image. In practise, this requires a sequence of steps including thinning the input image; running a search algorithm to find the top 5 parses; approximate type-level variability around each parse; and re-optimize the token-level variables. Extensive experimentation verify the capacity of the model and show that both classification and synthesis capabilities are on par with human performance. In class discussion, there were some concerns regarding the validity of the synthesis experiments. It seems that the characters in the data-set were generated by non-native writers, and entered using a computer mouse rather than a pen. Therefore, the general character appearances and stroke structure may be quite different from, and exhibit less variation than, characters generated by native writers using a standard writing device. Further, all human drawers are treated equally, when there may in fact be significant personal variation, in particular among native writers. Humans may not be particularly good at judging handwriting they are not familiar with, in particular since the human characters in the data-set, as mentioned above, may be exhibiting less variation than actual human hand writing. This could contribute to making it easier for the machine generated characters to pass the Turing test. It would be interesting to know whether a neural network, trained on this task, would be more suitable than humans to serve as judges in this situation. Another concern was with regards to method generality. Although a modified HBPL framework has been used for speech regocnition [2], several aspects of the method, e. In conclusion, the HBPL and related probabilistic programming methods offer a compelling alternative to the current domiant recognition paradigm of deep neural networks. Lake, Chia-ying Lee, James R.

3: Human-level concept learning through probabilistic program induction | Hacker News

title = "Human-level concept learning through probabilistic program induction", abstract = "People learning new concepts can often generalize successfully from just a single example, yet machine learning algorithms typically require tens or hundreds of examples to perform with similar accuracy.

HUMAN-LEVEL CONCEPT LEARNING THROUGH PROBABILISTIC PROGRAM INDUCTION pdf

4: Human-level concept learning through probabilistic program induction â€” NYU Scholars

Bayesian Program Learning is probably the most powerful of the technologies that come with the use of fully probabilistic programming languages. This is a subset of a technique known as "probabilistic generative modelling", in which one designs a model based on inference algorithms like MCMC or Particle Cascade, that generates samples based on.

5: Josh Tenenbaum's home page

Human-level concept learning through probabilistic program induction -- visual Turing tests By Brenden Lake, Ruslan Salakhutdinov, and Joshua Tenenbaum.

6: When machines learn like humans | Kurzweil

Probabilistic programs could capture these richer aspects of concept learning and use, but only with more abstract and complex structure than the programs studied here.

7: Human-level concept learning through probabilistic program induction | Read by QxMD

People learning new concepts can often generalize successfully from just a single example, yet machine learning algorithms typically require tens or hundreds of examples to perform with similar.

8: Human-level concept learning through probabilistic program induction | theberkeleyview

Human-level concept learning through probabilistic program induction - Papers We Love SG Building machines that see, learn, and think like people Concepts and Questions as Programs.

9: Visual Turing Test Demos

Human-level concept learning through probabilistic The best thing about this work is that by describing characters as stroke programs, you should be able to.

HUMAN-LEVEL CONCEPT LEARNING THROUGH PROBABILISTIC PROGRAM INDUCTION pdf

Engineering mechanics statics equilibrium Alexander I and Russian Reformism at the Beginning of the Nineteenth Century (Russian Studies in the Huma Making work pay in Nicaragua Russian Energy Policy During President Putins Tenure Diffusion and osmosis worksheet Introduction to insect study in Africa Dietary management of hepatitis The book of Letchworth Works by Sylvia Plath Iqbal poetry books Practical steam and hot water heating and ventilation Cleft Lip and or Palate The dead killing the living. Seeking convergence in policy and practice Japans security relations with China since 1989 Away in a Manger (Glad Tidings Board Books (Christmas Board Books) Spiritual Community Section I. The Landscape of K-12 Giving A Richard Lee Colvin Jay P. Greene The Leslie Lenkowsky Tom Lovel The Life Of Elbridge Gerry With Contemporary Letters Phlebotomy essentials 5th edition Lifes Obstacle Course: A Users Guide Democracy in california janiskee and masugi 4th edition JFK wants to know Juliette Low and the Girl Scouts A consequence of the popes criterion Indias food problem. Legal research guide to television broadcasting and program syndication Kids Book Of Basketball Primary prevention of psychiatric disorders The ladybird book of handwriting De la sonorite art et technique flute Breast cancer journal articles Masterpieces of Religious Verse The Creation Records Story No Good Deed Goes Unpunished 3./tREVIEW OF LITERATURE/t 15 Student and chief The Miracle of Recovery Cybercrime Cyberterrorism Cyberwarfare Innovation in industry