

481 **A Appendix**

482 **A.1 Relevance to NeurIPS**

483 The ML community is deeply motivated by a desire to have a positive impact on the world. This
484 desire is reflected in recent efforts in the ML community, such as NeurIPS’s requirement for the
485 inclusion of broader impacts statements for all submitted papers in 2020, the Resistance AI Workshop
486 at NeurIPS 2020 which investigated how AI concentrates power, and the Navigating the Broader
487 Impacts of AI Research at NeurIPS 2020 which sought to understand the impacts of ML research as a
488 whole on society. Understanding what the social impact of a paper, let alone the discipline, is difficult.
489 Merely looking at various benchmarks or broader impact statements, for example, is insufficient.
490 This paper attempts to begin to bridge this gap by seeking to understand the value commitments in
491 papers published at NeurIPS and a closely related conference, ICML. As such, this paper is highly
492 relevant to the NeurIPS audience. While research into core technical ML topics – reinforcement
493 learning, deep learning, optimization, etc. – are vital to NeurIPS and the wider ML community, so is
494 research on where these research areas stand with regard to societal impact, both in a positive and
495 negative manner, as well as the benefits they bring and to whom.

496 **A.2 Additional Methodological Details**

497 **A.2.1 Data Sources**

498 To determine the most-cited papers from each conference, we rely on the publicly-available Semantic
499 Scholar database, which includes bibliographic information for scientific papers, including citation
500 counts.⁶ Using this data, we chose the most cited papers from each of 2008, 2009, 2018, 2019
501 published at NeurIPS and ICML.

502 Like all bibliographic databases, Semantic Scholar is imperfect, and thus our selection includes
503 one paper that was actually published in 2010, and one that was retracted from NeurIPS prior to
504 publication (see §A.8 for details). In addition, the citations counts used to determine the most cited
505 papers reflect a static moment in time, and may differ from other sources.

506 Because all data used for this paper (aside from the actual annotations, which we contribute) have
507 been previously published at NeurIPS or ICML, we chose not to seek permission to annotate this
508 data from the original authors. Similarly, although it is possible that the original papers may contain
509 personally identifying information or offensive content, we rely on the fact that the original authors
510 contributed their work to the same community to which our own work is directed, and we thus believe
511 that the potential harm from this is minimal.

512 **A.2.2 Defining elite university**

513 To determine the list of elite universities, we follow Ahmed and Wahed [4], and rely on the QS World
514 University Rankings for the discipline of computer science. For 2018/19, we take the top 50 schools
515 from the CS rankings for 2018. For 2008/09, we take the top 50 schools from the CS rankings for
516 2011, as the closest year for which data is available.

517 **A.2.3 Defining big tech**

518 We used Abdalla and Abdalla’s [2] criterion to what is considered "big tech", which is comprised of:
519 Alibaba, Amazon, Apple, Element AI, Facebook, Google, Huawei, IBM, Intel, Microsoft, Nvidia,
520 Open AI, Samsung, and Uber. Furthermore, we added DeepMind to this list. We considered all other
521 companies as "non-big tech."

522 **A.3 Annotations**

523 We include the annotations of all papers in the supplementary zip file. To present a birds-eye view of
524 the value annotations, we present randomly selected examples of annotated sentences in section §A.7.
525 In addition, we present the frequency of occurrence for all values (prior to grouping) in Figure A.1
526 below.

⁶<http://s2-public-api.prod.s2.allenai.org/corpus/>

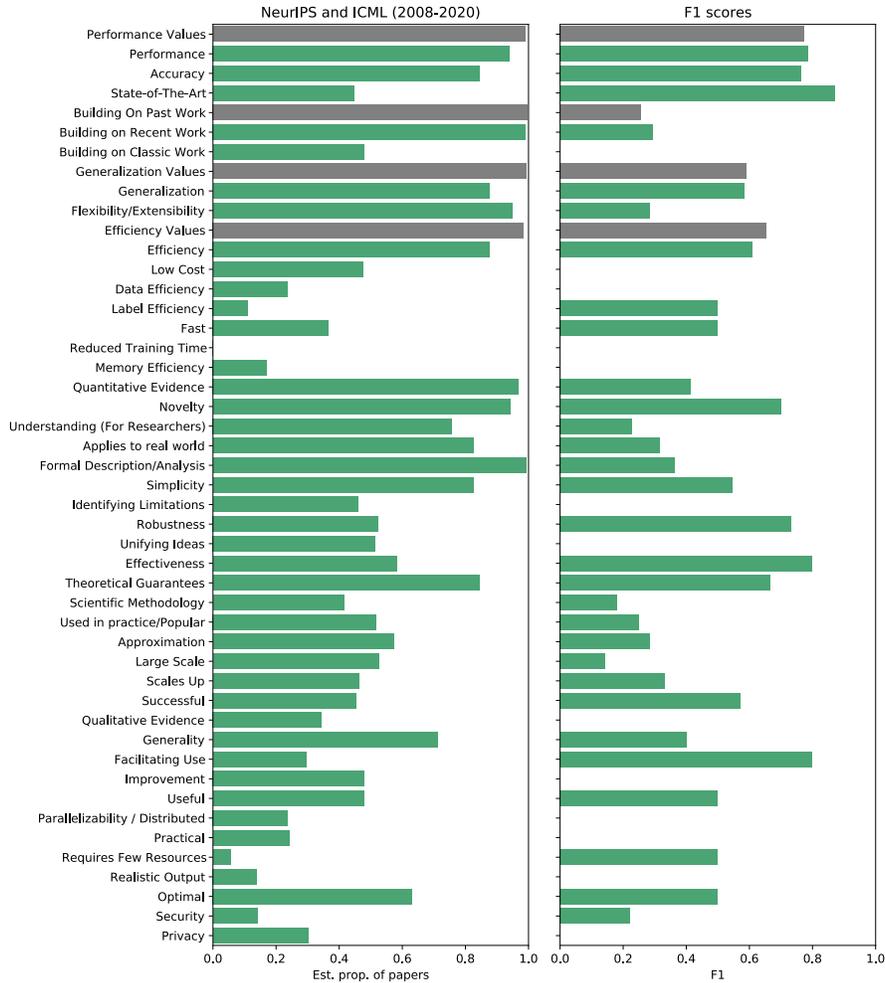


Figure A.2: Proportion of papers in from 2008–2020 (combining NeurIPS and ICML) predicted to have at least one sentence expressing each value (left), and estimated performance (F1) of the corresponding classifiers (right). Note that the overall performance of most classifiers is generally poor, indicating that the estimates on the left should be treated as unreliable in most cases. Grey bars represent the grouped values. Classifier were not trained for values with less than 20 representative sentences.

564 We then apply the classifiers trained above to each sentence in each paper. For each value, we then
 565 compute the proportion of papers (combining NeurIPS and ICML for this entire time period) that had
 566 at least one sentence predicted to exhibit that value. The overall proportions are shown in Figure A.2
 567 (left). As can be seen, the relative prevalence of values is broadly similar to our annotated sample,
 568 though many are predicted to occur with greater frequency, as expected. However, to reiterate, we
 569 should be highly skeptical of these findings, given the poor performance of the classifiers.

570 Finally, as an additional exploration, we focus on the Performance-related values (*Performance*,
 571 *Accuracy*, and *State-of-the-art*), which represent the overall most prevalent group in our annotations,
 572 and were relatively easy to identify using classification, and plot the estimated frequency over time
 573 for both conferences (For NeurIPS, which has better archival practices, we extend the analysis back to
 574 1987). We should again treat these results with caution, given all the caveats above, as well as the fact
 575 that we are now applying these classifiers outside the temporal range from which the annotations were
 576 collected. Nevertheless, the results, shown in Figure A.3, suggest that these values have gradually
 577 become more common in NeurIPS over time, reinforcing the contingent nature of the dominance of
 578 the current set of values. Further investigation is required, however, in order to verify this finding.

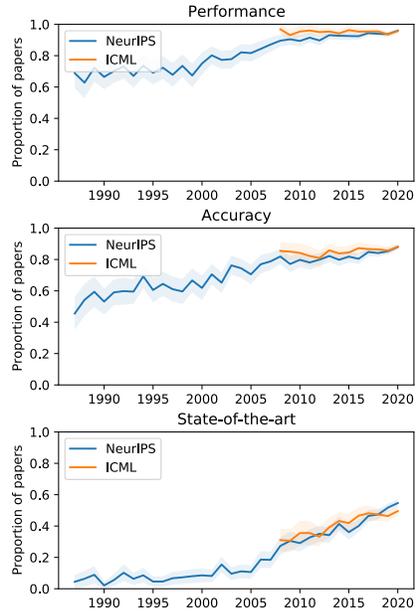


Figure A.3: Proportion of papers per year (of those published in ICML and NeurIPS) that are classified as having at least one sentence expressing *Performance*, *Accuracy*, or *State-of-the-art*, (top, middle, and bottom), based on simple text classifiers trained on our annotations. Bands show ± 2 standard deviations, reflecting the changing overall number of papers per year.

579 A.5 Code and Reproducibility

580 We include, in the supplementary zip file, the code used for all data analysis; in particular, we include
 581 the code used to run the text classification experiments and generate all figures in the paper. The text
 582 classification experiments were run on a 2019 Macbook Air.

583 A.6 Potential Negative Societal Impacts

584 Because this paper relies only on manual annotation of papers already published at NeurIPS and
 585 ICML, we believe that the potential negative societal impacts of carrying out these annotations and
 586 sharing them are minimal. However, we still briefly comment on this here.

587 First, in terms of the specific concerns highlighted in the NeurIPS call for papers, we believe our
 588 annotation work poses no risk to living beings, human rights concerns, threats to livelihoods, etc.
 589 Similarly, all annotators are co-authors on this paper, thus there was no risk to participants, beyond
 590 what we chose to take on for ourselves.

591 One area of potential concern might be in terms of unintentionally casting certain authors in a negative
 592 light, or unintentionally contributing to harmful tensions within the ML community. In order to
 593 minimize the risk of the former, we have chosen to include randomly selected examples but omit
 594 author attributions from quoted sources in the main paper. However, we do include a full list of
 595 cited papers below, so as to both acknowledge this work, but not draw attention to any one particular
 596 source.

597 Although our intention is to broaden the conversation, we do acknowledge that some authors may
 598 perceive our work as being not representative of the type of work they would like to see at NeurIPS,
 599 and possibly detrimental to the conference. However, because of the prominence of machine learning
 600 today, we feel it is especially important to have these conversations at the premier venues, and hope
 601 that our paper will be the basis for useful conversations and future work.

602 A.7 Random Examples

603 The list below contains 100 random examples drawn from the annotated data, along with the set of
604 annotated values for each. These sentences were annotated for values within the context of the entire
605 paper.

- 606 • The problem of minimizing the rank of a matrix variable subject to certain constraints arises
607 in many fields including machine learning, automatic control, and image compression. **Used**
608 **in practice/Popular**
- 609 • Locality-sensitive hashing [6] is an effective technique that performs approximate nearest
610 neighbor searches in time that is sub-linear in the size of the database **Approximation,**
611 **Building on recent work, Effectiveness, Fast**
- 612 • In the finite case, analysis of optimization and generalization of fully-trained nets is of course
613 an open problem **Formal description/analysis, Generalization**
- 614 • So to achieve adversarial robustness, a classifier must generalize in a stronger sense. **Gener-**
615 **alization, Robustness**
- 616 • Robustness to label corruption is similarly improved by wide margins, such that pre-training
617 alone outperforms certain task-specific methods, sometimes even after combining these
618 methods with pre-training. **Performance, Robustness, Understanding (for researchers)**
- 619 • RBMs have been particularly successful in classification problems either as feature extractors
620 for text and image data (Gehler et al., 2006) or as a good initial training phase for deep neural
621 network classifiers (Hinton, 2007). **Building on recent work, Flexibility/Extensibility,**
622 **Successful**
- 623 • Our theoretical analysis naturally leads to a new formulation of adversarial defense which
624 has several appealing properties; in particular, it inherits the benefits of scalability to large
625 datasets exhibited by Tiny ImageNet, and the algorithm achieves state-of-the-art performance
626 on a range of benchmarks while providing theoretical guarantees. **Robustness, Scales up,**
627 **Security, Theoretical guarantees**
- 628 • The current paper focuses on the training loss, but does not address the test loss. **General-**
629 **ization**
- 630 • This result is significant since stochastic methods are highly preferred for their efficiency
631 over deterministic gradient methods in machine learning applications. **Efficiency**
- 632 • Ranking, which is to sort objects based on certain factors, is the central problem of
633 applications such as information retrieval (IR) and information filtering. **Applies to real**
634 **world, Used in practice/Popular**
- 635 • This subspace is important, because, when projected onto this subspace, the means of the
636 distributions are well-separated, yet the typical distance between points from the same
637 distribution is smaller than in the original space. **Important**
- 638 • Overall, the existence of such adversarial examples raises concerns about the robustness of
639 current classifiers. **Identifying limitations, Robustness**
- 640 • We have shown that biased compressors if naively used can lead to bad generalization, and
641 even non-convergence. **Formal description/analysis, Generalization**
- 642 • Bartlett and Mendelson [2002] provide a generalization bound for Lipschitz loss functions.
643 **Building on classic work, Generalization**
- 644 • The principal advantage of taking this “lateral” approach arises from the fact that compact
645 representation in trajectory space is better motivated physically than compact representation
646 in shape space **Realistic world model**
- 647 • In this paper, we show that gradient descent on deep overparametrized networks can obtain
648 zero training loss **Formal description/analysis, Theoretical guarantees**
- 649 • Moreover, web queries often have different meanings for different users (a canonical example
650 is the query jaguar) suggesting that a ranking with diverse documents may be preferable.
651 **Diverse output, User influence**

- 652 • We include human performance estimates for all benchmark tasks, which verify that sub-
653 stantial headroom exists between a strong BERT-based baseline and human performance.
654 **Learning from humans, Performance**
- 655 • In this paper we propose a simple and fast algorithm SVP (Singular Value Projection) for rank
656 minimization under affine constraints (ARMP) and show that SVP recovers the minimum
657 rank solution for affine constraints that satisfy a restricted isometry property (RIP). **Fast,**
658 **Novelty, Simplicity**
- 659 • We use standard formalization of multiclass classification, where data consists of sample x
660 and its label y (an integer from 1 to k). **Building on classic work**
- 661 • A number of recent works has shown that the low rank solution can be recovered exactly via
662 minimizing the trace norm under certain conditions (Recht et al., 2008a; Recht et al., 2008b;
663 Candes Recht, 2008). **Building on recent work**
- 664 • This difficulty has necessitated the use of a heuristic inference procedure, that nonetheless
665 was accurate enough for successful learning. **Accuracy, Successful**
- 666 • We illustrate such potential by measuring search space properties relevant to architecture
667 search. **Quantitative evidence (e.g. experiments)**
- 668 • Deep architectures consist of feature detector units arranged in layers. Lower layers detect
669 simple features and feed into higher layers, which in turn detect more complex features.
670 **Simplicity**
- 671 • This makes the updates hard to massively parallelize at a coarse, data-parallel level (e.g., by
672 computing the updates in parallel and summing them together centrally) without losing the
673 critical stochastic nature of the updates. **Large scale, Parallelizability / distributed**
- 674 • This suggests future work on model robustness should evaluate proposed methods with
675 pretraining in order to correctly gauge their utility, and some work could specialize pre-
676 training for these downstream tasks. **Robustness**
- 677 • Adversarial training remains among the most trusted defenses, but it is nearly intractable on
678 large-scale problems. **Scales up, Security**
- 679 • For complex robots such as humanoids or light-weight arms, it is often hard to model the
680 system sufficiently well and, thus, modern regression methods offer a viable alternative
681 [7,8]. **Realistic world model**
- 682 • In contrast to prior work that operates in this goal-setting model, we use states as goals
683 directly, which allows for simple and fast training of the lower layer. **Reduced training**
684 **time, Simplicity**
- 685 • Meanwhile, using less resources tends to produce less compelling results (Negrinho Gordon,
686 2017; Baker et al., 2017a). **Requires few resources**
- 687 • This finding represents an exciting opportunity for defense against neural fake news: the
688 best models for generating neural disinformation are also the best models at detecting it.
689 **Applies to real world**
- 690 • Our strong empirical results suggest that randomized smoothing is a promising direction
691 for future research into adversarially robust classification. **Quantitative evidence (e.g.**
692 **experiments), Robustness, Security**
- 693 • We then turn our attention to identifying the roots of BatchNorm’s success. **Successful,**
694 **Understanding (for researchers)**
- 695 • We also report the results of large-scale experiments comparing these three methods which
696 demonstrate the benefits of the mixture weight method: this method consumes less resources,
697 while achieving a performance comparable to that of standard approaches. **Large scale,**
698 **Performance, Requires few resources**
- 699 • This paper does not cover the generalization of over-parameterized neural networks to
700 the test data. **Avoiding train/test discrepancy, Generalization**
- 701 • While there has been success with robust classifiers on simple datasets [31, 36, 44, 48],
702 more complicated datasets still exhibit a large gap between “standard” and robust accuracy
703 [3, 11]. **Applies to real world, Robustness, Successful**

- 704 • In this paper, we have shown theoretically how independence between examples can make
705 the actual effect much smaller. **Novelty, Theoretical guarantees**
- 706 • We provide empirical evidence that several recently-used methods for estimating the proba-
707 bility of held-out documents are inaccurate and can change the results of model comparison.
708 **Accuracy, Building on recent work, Quantitative evidence (e.g. experiments)**
- 709 • This agreement is robust across different architectures, optimization methods, and loss
710 functions **Robustness**
- 711 • Unfortunately, due to the slow-changing policy in an actor-critic setting, the current and
712 target value estimates remain too similar to avoid maximization bias. **Accuracy**
- 713 • As a future work, we are pursuing a better understanding of probabilistic distributions on
714 the Grassmann manifold. **Understanding (for researchers)**
- 715 • We also view these results as an opportunity to encourage the community to pursue a more
716 systematic investigation of the algorithmic toolkit of deep learning and the underpinnings of
717 its effectiveness. **Effectiveness, Understanding (for researchers)**
- 718 • This challenge is further exacerbated in continuous state and action spaces, where a separate
719 actor network is often used to perform the maximization in Q-learning. **Performance**
- 720 • The vulnerability of neural networks to adversarial perturbations has recently been a source
721 of much discussion and is still poorly understood. **Robustness, Understanding (for re-
722 searchers)**
- 723 • Most of the evaluation methods described in this paper extend readily to more complicated
724 topic models— including non-parametric versions based on hierarchical Dirichlet processes
725 (Teh et al., 2006)—since they only require a MCMC algorithm for sampling the latent
726 topic assignments z for each document and a way of evaluating probability $P(w | z, \theta, m)$.
727 **Flexibility/Extensibility, Understanding (for researchers)**
- 728 • In a formulation closely related to the dual problem, we have: $w^* = \operatorname{argmin}_w F(w) + c \sum_{i=1}^n \langle w, x_i \rangle - \sum_{i=1}^n \langle w, y_i \rangle$ (2) where, instead of regularizing, a hard restriction over the parameter
729 space is imposed (by the constant c). **Formal description/analysis**
- 730 • Second, we evaluate a surrogate loss function from four aspects: (a) consistency, (b)
731 soundness, (c) mathematical properties of continuity, differentiability, and con-
732 vexity, and (d) computational efficiency in learning. **Efficiency**
- 733 • This leads to two natural questions that we try to answer in this paper: (1) Is it feasible to
734 perform optimization in this very large feature space with cost which is polynomial in the
735 size of the input space? **Performance**
- 736 • Despite its pervasiveness, the exact reasons for BatchNorm’s effectiveness are still poorly
737 understood. **Understanding (for researchers)**
- 738 • We have presented confidenceweighted linear classifiers, a new learning method designed
739 for NLP problems based on the notion of parameter confidence. **Novelty**
- 740 • In addition, the experiments reported here suggest that (like other strategies recently proposed
741 to train deep deterministic or stochastic neural networks) the curriculum strategies appear
742 on the surface to operate like a regularizer, i.e., their beneficial effect is most pronounced on
743 the test set. **Beneficence, Quantitative evidence (e.g. experiments)**
- 744 • These give further inside into hash-spaces and explain previously made empirical observa-
745 tions. **Understanding (for researchers)**
- 746 • This means that current algorithms reach their limit at problems of size 1TB whenever the
747 algorithm is I/O bound (this amounts to a training time of 3 hours), or even smaller problems
748 whenever the model parametrization makes the algorithm CPU bound. **Memory efficiency,
749 Reduced training time**
- 750 • Much of the results presented were based on the assumption that the target distribution is
751 some mixture of the source distributions. **Valid assumptions**
- 752 • Empirical investiga- tion revealed that this agrees well with actual training dynamics and
753 predictive distributions across fully-connected, convolutional, and even wide residual net-
754 work architectures, as well as with different optimizers (gradient descent, momentum,
755 mini-batching) and loss functions (MSE, cross-entropy). **Generalization, Quantitative
756 evidence (e.g. experiments), Understanding (for researchers)**
- 757

- 758 • We design a new spectral norm that encodes this a priori assumption, without the prior
759 knowledge of the partition of tasks into groups, resulting in a new convex optimization
760 formulation for multi-task learning. **Novelty**
- 761 • Recent progress in natural language generation has raised dual-use concerns. **Progress**
- 762 • These kernel functions can be used in shallow architectures, such as support vector machines
763 (SVMs), or in deep kernel-based architectures that we call multilayer kernel machines
764 (MKMs). **Flexibility/Extensibility**
- 765 • Using MCMC instead of variational methods for approximate inference in Bayesian matrix
766 factorization models leads to much larger improvements over the MAP trained models,
767 which suggests that the assumptions made by the variational methods about the structure of
768 the posterior are not entirely reasonable. **Understanding (for researchers)**
- 769 • In particular, the deep belief network (DBN) (Hinton et al., 2006) is a multilayer generative
770 model where each layer encodes statistical dependencies among the units in the layer
771 below it; it is trained to (approximately) maximize the likelihood of its training data.
772 **Approximation, Data efficiency**
- 773 • Furthermore, the learning accuracy and performance of our LGP approach will be compared
774 with other important standard methods in Section 4, e.g., LWPR [8], standard GPR [1],
775 sparse online Gaussian process regression (OGP) [5] and -support vector regression (-SVR)
776 [11], respectively **Accuracy, Performance, Quantitative evidence (e.g. experiments)**
- 777 • • propose a simple method based on weighted minibatches to stochastically train with
778 arbitrary weights on the terms of our decomposition without any additional hyperparameters.
779 **Efficiency, Simplicity**
- 780 • For example, Ng (2004) examined the task of PAC learning a sparse predictor and analyzed
781 cases in which an 1 constraint results in better solutions than an 2 constraint. **Building on**
782 **recent work**
- 783 • Graph Convolutional Networks (GCNs) (Kipf Welling, 2017) are an efficient variant of
784 Convolutional Neural Networks (CNNs) on graphs. GCNs stack layers of learned first-order
785 spectral filters followed by a nonlinear activation function to learn graph representations.
786 **Efficiency**
- 787 • This is a linear convergence rate. **Building on recent work, Efficiency, Quantitative**
788 **evidence (e.g. experiments), Theoretical guarantees**
- 789 • However, as we observe more interactions, this could emerge as a clear feature. **Building**
790 **on recent work, Data efficiency**
- 791 • Here we propose the first method that supports arbitrary low accuracy and even biased
792 compression operators, such as in (Alistarh et al., 2018; Lin et al., 2018; Stich et al., 2018).
793 **Accuracy, Novelty**
- 794 • Much recent work has been done on understanding under what conditions we can learn a
795 mixture model. **Understanding (for researchers)**
- 796 • For this reason, we present an extension of the standard greedy OMP algorithm that can be
797 applied to general struc- tured sparsity problems, and more importantly, meaningful sparse
798 recovery bounds can be obtained for this algorithm. **Building on recent work**
- 799 • In this paper we show that this assumption is indeed nec-essary: by considering a simple
800 yet prototypical example of GAN training we analytically show that (unregularized) GAN
801 training is not always locally convergent **Formal description/analysis**
- 802 • Overestimation bias is a property of Q-learning in which the maximization of a noisy value
803 estimate induces a consistent overestimation **Accuracy**
- 804 • This drawback prevents GPR from applications which need large amounts of training
805 data and require fast computation, e.g., online learning of inverse dynamics model for
806 model-based robot control **Fast, Large scale**
- 807 • This is problematic since we find there are techniques which do not comport well with
808 pre-training; thus some evaluations of robustness are less representative of real-world
809 performance than previously thought. **Applies to real world, Performance, Robustness**

- 810 • Approximation of this prior structure through simple, efficient hyperparameter optimization steps is sufficient to achieve these performance gains **Approximation, Efficiency, Performance, Simplicity**
- 811
- 812
- 813 • The second mysterious phenomenon in training deep neural networks is “deeper networks are harder to train.” **Performance**
- 814
- 815 • However, the definition of our metric is sufficiently general that it could easily be used to test, for example, invariance of auditory features to rate of speech, or invariance of textual features to author identity. **Generalization**
- 816
- 817
- 818 • In Sec. 6 we test the proposed algorithm for face recognition and object categorization tasks. **Applies to real world, Quantitative evidence (e.g. experiments)**
- 819
- 820 • It is possible to train classification RBMs directly for classification performance; the gradient is fairly simple and certainly tractable. **Performance**
- 821
- 822 • Figure 1 contrasts these two approaches. Defining and evaluating models using ODE solvers has several benefits: **Beneficence**
- 823
- 824 • They claim to achieve $12 \cdot 2$ radius of 3 (for images with pixels in $[0, 1]$). **Generalization, Robustness**
- 825
- 826 • Two commonly used penalties are the 1- norm and the square of the 2-norm of w . **Used in practice/Popular**
- 827
- 828 • What should platforms do? Video-sharing platforms like YouTube use deep neural networks to scan videos while they are uploaded, to filter out content like pornography (Hosseini et al., 2017). **Applies to real world**
- 829
- 830
- 831 • We mention various properties of this penalty, and provide conditions for the consistency of support estimation in the regression setting. Finally, we report promising results on both simulated and real data **Applies to real world**
- 832
- 833
- 834 • There could be a separate feature for “high school student,” “male,” “athlete,” and “musician” and the presence or absence of each of these features is what defines each person and determines their relationships. **Building on recent work**
- 835
- 836
- 837 • So, the over-parameterized convergence theory of DNN is much simpler than that of RNN. **Simplicity, Understanding (for researchers)**
- 838
- 839 • Other threat models are possible: for instance, an adversary might generate comments or have entire dialogue agents, they might start with a human-written news article and modify a few sentences, and they might fabricate images or video. **Learning from humans**
- 840
- 841
- 842 • More generally, we hope that future work will be able to avoid relying on obfuscated gradients (and other methods that only prevent gradient descent-based attacks) for perceived robustness, and use our evaluation approach to detect when this occurs. **Generality, Robustness**
- 843
- 844
- 845
- 846 • For example, the learned linear combination does not consistently outperform either the uniform combination of base kernels or simply the best single base kernel (see, for example, UCI dataset experiments in [9, 12], see also NIPS 2008 workshop). **Performance**
- 847
- 848
- 849 • Our main contributions are: • We analyze GP-UCB, an intuitive algorithm for GP optimization, when the function is either sampled from a known GP, or has low RKHS norm. **Optimal**
- 850
- 851
- 852 • For the standard linear setting, Dani et al. (2008) provide a near-complete characterization explicitly dependent on the dimensionality. In the GP setting, the challenge is to characterize complexity in a different manner, through properties of the kernel function. **Building on classic work**
- 853
- 854
- 855
- 856 • This allows us to map each architecture A to its approximate hyperparameteroptimized accuracy **Accuracy**
- 857
- 858 • Unfortunately, they could only apply their method to linear networks. **Flexibility/Extensibility**
- 859
- 860 • The strength of the adversary then allows for a trade-off between the enforced prior, and the data-dependent features. **Understanding (for researchers)**
- 861

- 862 • We observe that the computational bottleneck of NAS is the training of each child model
863 to convergence, only to measure its accuracy whilst throwing away all the trained weights.
864 **Accuracy**
- 865 • We show that the number of subproblems need only be logarithmic in the total number of
866 possible labels, making this approach radically more efficient than others. **Efficiency**
- 867 • We establish a new notion of quadratic approximation of the neural network, and connect
868 it to the SGD theory of escaping saddle points. **Novelty, Unifying ideas or integrating**
869 **components**
- 870 • In this work, we decompose the prediction error for adversarial examples (robust error) as
871 the sum of the natural (classification) error and boundary error, and provide a differentiable
872 upper bound using the theory of classification-calibrated loss, which is shown to be the
873 tightest possible upper bound uniform over all probability distributions and measurable
874 predictors. **Accuracy, Robustness, Theoretical guarantees**
- 875 • A limit on the number of queries can be a result of limits on other resources, such as a time
876 limit if inference time is a bottleneck or a monetary limit if the attacker incurs a cost for
877 each query. **Applies to real world, Low cost, Requires few resources**
- 878 • Preliminary experiments demonstrate that it is significantly faster than batch alternatives on
879 large datasets that may contain millions of training examples, yet it does not require learning
880 rate tuning like regular stochastic gradient descent methods. **Quantitative evidence (e.g.**
881 **experiments), Reduced training time**
- 882 • SuperGLUE is available at super.gluebenchmark.com. **Facilitating use (e.g. sharing code)**

883 A.8 Full List of Cited Papers

884 The full list of annotated papers is given below, along with the annotated scores (in square brackets)
885 for *Discussion of Negative Potential* [left] (0 = Doesn't mention negative potential; 1 = Mentions
886 but does not discuss negative potential; 2 = Discusses negative potential) and *Justification* [right] (0
887 = Doesn't rigorously justify how it achieves technical goal; 1 = Justifies how it achieves technical
888 goal but no mention of societal need; 2 = States but does not justify how it connects to a societal
889 need; 3 = States and somewhat justifies how it connects to a societal need; 4 = States and rigorously
890 justifies how it connects to a societal need). Note that due to minor errors in the data sources used,
891 the distribution of papers over venues and years is not perfectly balanced. For the same reason, the
892 list also contains one paper from 2010 (rather than 2009), as well as one paper that was retracted
893 before publication at NeurIPS (marked with a *).

- 894 • Mingxing Tan, Quoc Le. [EfficientNet: Rethinking Model Scaling for Convolutional Neural
895 Networks](#). In *Proceedings of ICML*, 2019. [0/1]
- 896 • Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, Ruosong Wang. [Fine-Grained Analysis of
897 Optimization and Generalization for Overparameterized Two-Layer Neural Networks](#). In
898 *Proceedings of ICML*, 2019. [0/1]
- 899 • Jeremy Cohen, Elan Rosenfeld, Zico Kolter. [Certified Adversarial Robustness via Random-
900 ized Smoothing](#). In *Proceedings of ICML*, 2019. [0/1]
- 901 • Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, Michael Jordan.
902 [Theoretically Principled Trade-off between Robustness and Accuracy](#). In *Proceedings of
903 ICML*, 2019. [0/2]
- 904 • Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, Tie-Yan Liu. [MASS: Masked Sequence to
905 Sequence Pre-training for Language Generation](#). In *Proceedings of ICML*, 2019. [0/1]
- 906 • Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, Kilian Weinberger.
907 [Simplifying Graph Convolutional Networks](#). In *Proceedings of ICML*, 2019. [0/1]
- 908 • Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, Vaishal Shankar. [Do ImageNet
909 Classifiers Generalize to ImageNet?](#) In *Proceedings of ICML*, 2019. [0/2]
- 910 • Justin Gilmer, Nicolas Ford, Nicholas Carlini, Ekin Cubuk. [Adversarial Examples Are a
911 Natural Consequence of Test Error in Noise](#). In *Proceedings of ICML*, 2019. [0/1]

- 912 • Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, Frank Hutter.
913 [NAS-Bench-101: Towards Reproducible Neural Architecture Search.](#) In *Proceedings of*
914 *ICML, 2019.* [0/2]
- 915 • Dan Hendrycks, Kimin Lee, Mantas Mazeika. [Using Pre-Training Can Improve Model](#)
916 [Robustness and Uncertainty.](#) In *Proceedings of ICML, 2019.* [0/1]
- 917 • Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian Stich, Martin Jaggi. [Error Feedback](#)
918 [Fixes SignSGD and other Gradient Compression Schemes.](#) In *Proceedings of ICML, 2019.*
919 [0/1]
- 920 • Anastasia Koloskova, Sebastian Stich, Martin Jaggi. [Decentralized Stochastic Optimization](#)
921 [and Gossip Algorithms with Compressed Communication.](#) In *Proceedings of ICML, 2019.*
922 [0/2]
- 923 • Han Zhang, Ian Goodfellow, Dimitris Metaxas, Augustus Odena. [Self-Attention Generative](#)
924 [Adversarial Networks.](#) In *Proceedings of ICML, 2019.* [0/1]
- 925 • Zeyuan Allen-Zhu, Yuanzhi Li, Zhao Song. [A Convergence Theory for Deep Learning via](#)
926 [Over-Parameterization.](#) In *Proceedings of ICML, 2019.* [0/1]
- 927 • Simon Du, Jason Lee, Haochuan Li, Liwei Wang, Xiyu Zhai. [Gradient Descent Finds Global](#)
928 [Minima of Deep Neural Networks.](#) In *Proceedings of ICML, 2019.* [0/1]
- 929 • Anish Athalye, Nicholas Carlini, David Wagner. [Obfuscated Gradients Give a False Sense](#)
930 [of Security: Circumventing Defenses to Adversarial Examples.](#) In *Proceedings of ICML,*
931 *2018.* [0/2]
- 932 • Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, Jeff Dean. [Efficient Neural Architecture](#)
933 [Search via Parameters Sharing.](#) In *Proceedings of ICML, 2018.* [0/1]
- 934 • Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, Sergey Levine. [Soft Actor-Critic: Off-Policy](#)
935 [Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor.](#) In *Proceedings*
936 *of ICML, 2018.* [0/2]
- 937 • Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam
938 Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, Koray Kavukcuoglu. [IMPALA:](#)
939 [Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures.](#) In
940 *Proceedings of ICML, 2018.* [0/1]
- 941 • Scott Fujimoto, Herke Hoof, David Meger. [Addressing Function Approximation Error in](#)
942 [Actor-Critic Methods.](#) In *Proceedings of ICML, 2018.* [0/1]
- 943 • Hyunjik Kim, Andriy Mnih. [Disentangling by Factorising.](#) In *Proceedings of ICML, 2018.*
944 [0/0]
- 945 • Lars Mescheder, Andreas Geiger, Sebastian Nowozin. [Which Training Methods for GANs](#)
946 [do actually Converge?](#) In *Proceedings of ICML, 2018.* [0/1]
- 947 • Sanjeev Arora, Rong Ge, Behnam Neyshabur, Yi Zhang. [Stronger generalization bounds for](#)
948 [deep nets via a compression approach.](#) In *Proceedings of ICML, 2018.* [0/3]
- 949 • Andrew Ilyas, Logan Engstrom, Anish Athalye, Jessy Lin. [Black-box Adversarial Attacks](#)
950 [with Limited Queries and Information.](#) In *Proceedings of ICML, 2018.* [0/2]
- 951 • Niranjan Srinivas, Andreas Krause, Sham Kakade, Matthias Seeger. [Gaussian Process](#)
952 [Optimization in the Bandit Setting: No Regret and Experimental Design.](#) In *Proceedings of*
953 *ICML, 2010.* [0/1]
- 954 • Honglak Lee, Roger Grosse, Rajesh Ranganath and Andrew Ng. [Convolutional deep belief](#)
955 [networks for scalable unsupervised learning of hierarchical representations.](#) In *Proceedings*
956 *of ICML, 2009.* [0/1]
- 957 • Julien Mairal, Francis Bach, Jean Ponce and Guillermo Sapiro. [Online dictionary learning](#)
958 [for sparse coding.](#) In *Proceedings of ICML, 2009.* [0/1]
- 959 • Yoshua Bengio, Jerome Louradour, Ronan Collobert and Jason Weston. [Curriculum learning.](#)
960 In *Proceedings of ICML, 2009.* [0/1]
- 961 • Laurent Jacob, Guillaume Obozinski and Jean-Philippe Vert. [Group Lasso with Overlaps](#)
962 [and Graph Lasso.](#) In *Proceedings of ICML, 2009.* [0/3]

- 963 • Chun-Nam Yu and Thorsten Joachims. [Learning structural SVMs with latent variables.](#) In
964 *Proceedings of ICML, 2009.* [0/2]
- 965 • Kilian Weinberger, Anirban Dasgupta, Josh Attenberg, John Langford and Alex Smola.
966 [Feature hashing for large scale multitask learning.](#) In *Proceedings of ICML, 2009.* [0/2]
- 967 • Hanna Wallach, Iain Murray, Ruslan Salakhutdinov, and David Mimno. [Evaluation methods](#)
968 [for topic models.](#) In *Proceedings of ICML, 2009.* [0/1]
- 969 • Kamalika Chaudhuri, Sham Kakade, Karen Livescu and Karthik Sridharan. [Multi-view](#)
970 [clustering via canonical correlation analysis.](#) In *Proceedings of ICML, 2009.* [0/2]
- 971 • Shuiwang Ji and Jieping Ye. [An accelerated gradient method for trace norm minimization.](#)
972 In *Proceedings of ICML, 2009.* [0/3]
- 973 • Junzhou Huang, Tong Zhang and Dimitris Metaxas. [Learning with structured sparsity.](#) In
974 *Proceedings of ICML, 2009.* [0/1]
- 975 • Rajat Raina, Anand Madhavan and Andrew Ng. [Large-scale deep unsupervised learning](#)
976 [using graphics processors.](#) In *Proceedings of ICML, 2009.* [0/2]
- 977 • Ronan Collobert and Jason Weston. [A unified architecture for natural language processing:](#)
978 [deep neural networks with multitask learning.](#) In *Proceedings of ICML, 2008.* [0/2]
- 979 • Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. [Extracting](#)
980 [and composing robust features with denoising autoencoders.](#) In *Proceedings of ICML, 2008.*
981 [0/1]
- 982 • Ruslan Salakhutdinov and Andriy Mnih. [Bayesian probabilistic matrix factorization using](#)
983 [Markov chain Monte Carlo.](#) In *Proceedings of ICML, 2008.* [0/1]
- 984 • John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. [Efficient projections](#)
985 [onto the \$l_1\$ -ball for learning in high dimensions.](#) In *Proceedings of ICML, 2008.* [0/1]
- 986 • Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S. Sathya Keerthi, and S. Sundararajan. [A](#)
987 [dual coordinate descent method for large-scale linear SVM.](#) In *Proceedings of ICML, 2008.*
988 [0/1]
- 989 • Tijmen Tieleman. [Training restricted Boltzmann machines using approximations to the](#)
990 [likelihood gradient.](#) In *Proceedings of ICML, 2008.* [0/1]
- 991 • Hugo Larochelle and Yoshua Bengio. [Classification using discriminative restricted Boltz-](#)
992 [mann machines.](#) In *Proceedings of ICML, 2008.* [0/1]
- 993 • Jihun Hamm and Daniel Lee. [Grassmann discriminant analysis: a unifying view on subspace-](#)
994 [based learning.](#) In *Proceedings of ICML, 2008.* [0/1]
- 995 • Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. [Listwise Approach to](#)
996 [Learning to Rank - Theory and Algorithm.](#) In *Proceedings of ICML, 2008.* [0/1]
- 997 • Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. [Learning diverse rankings with](#)
998 [multi-armed bandits.](#) In *Proceedings of ICML, 2008.* [0/1]
- 999 • Mark Dredze, Koby Crammer, and Fernando Pereira. [Confidence-weighted linear classifica-](#)
1000 [tion.](#) In *Proceedings of ICML, 2008.* [0/1]
- 1001 • Ruslan Salakhutdinov and Iain Murray. [On the quantitative analysis of deep belief networks.](#)
1002 In *Proceedings of ICML, 2008.* [0/1]
- 1003 • Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, Quoc
1004 V. Le. [XLNet: Generalized Autoregressive Pretraining for Language Understanding.](#) In
1005 *Proceedings of NeurIPS, 2019.* [0/1]
- 1006 • Alexis CONNEAU, Guillaume Lample. [Cross-lingual Language Model Pretraining.](#) In
1007 *Proceedings of NeurIPS, 2019.* [0/4]
- 1008 • Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, Alek-
1009 sander Madry. [Adversarial Examples Are Not Bugs, They Are Features.](#) In *Proceedings of*
1010 *NeurIPS, 2019.* [0/1]
- 1011 • Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha
1012 Sohl-Dickstein, Jeffrey Pennington. [Wide Neural Networks of Any Depth Evolve as Linear](#)
1013 [Models Under Gradient Descent.](#) In *Proceedings of NeurIPS, 2019.* [0/1]

- 1014 • David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, Colin
1015 A. Raffel. [MixMatch: A Holistic Approach to Semi-Supervised Learning](#). In *Proceedings*
1016 *of NeurIPS*, 2019. [0/1]
- 1017 • Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
1018 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas
1019 Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,
1020 Benoit Steiner, Lu Fang, Junjie Bai, Soumith Chintala. [PyTorch: An Imperative Style](#).
1021 [High-Performance Deep Learning Library](#). In *Proceedings of NeurIPS*, 2019. [0/1]
- 1022 • Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, Russ R. Salakhutdinov, Ruosong Wang.
1023 [On Exact Computation with an Infinitely Wide Neural Net](#). In *Proceedings of NeurIPS*,
1024 2019. [0/1]
- 1025 • Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao,
1026 Ming Zhou, Hsiao-Wuen Hon. [Unified Language Model Pre-training for Natural Language](#)
1027 [Understanding and Generation](#). In *Proceedings of NeurIPS*, 2019. [0/1]
- 1028 • Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph
1029 Studer, Larry S. Davis, Gavin Taylor, Tom Goldstein. [Adversarial Training for Free!](#) In
1030 *Proceedings of NeurIPS*, 2019. [0/3]
- 1031 • Jiasen Lu, Dhruv Batra, Devi Parikh, Stefan Lee. [ViLBERT: Pretraining Task-Agnostic](#)
1032 [Visiolinguistic Representations for Vision-and-Language Tasks](#). In *Proceedings of NeurIPS*,
1033 2019. [0/1]
- 1034 • Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix
1035 Hill, Omer Levy, Samuel Bowman. [SuperGLUE: A Stickier Benchmark for General-Purpose](#)
1036 [Language Understanding Systems](#). In *Proceedings of NeurIPS*, 2019. [1/1]
- 1037 • Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska
1038 Roesner, Yejin Choi. [Defending Against Neural Fake News](#). In *Proceedings of NeurIPS*,
1039 2019. [2/4]
- 1040 • Yuan Cao, Quanquan Gu. [Generalization Bounds of Stochastic Gradient Descent for Wide](#)
1041 [and Deep Neural Networks](#). In *Proceedings of NeurIPS*, 2019. [0/1]
- 1042 • Florian Tramer, Dan Boneh. [Adversarial Training and Robustness for Multiple Perturbations](#).
1043 In *Proceedings of NeurIPS*, 2019. [0/2]
- 1044 • Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, Percy S. Liang. [Unla-](#)
1045 [beled Data Improves Adversarial Robustness](#). In *Proceedings of NeurIPS*, 2019. [0/1]
- 1046 • Lars Maaløe, Marco Fraccaro, Valentin Liévin, Ole Winther. [BIVA: A Very Deep Hierarchy](#)
1047 [of Latent Variables for Generative Modeling](#). In *Proceedings of NeurIPS*, 2019. [0/1]
- 1048 • Zeyuan Allen-Zhu, Yuanzhi Li, Yingyu Liang. [Learning and Generalization in Overparam-](#)
1049 [eterized Neural Networks, Going Beyond Two Layers](#). In *Proceedings of NeurIPS*, 2019.
1050 [0/1]
- 1051 • Durk P. Kingma, Prafulla Dhariwal. [Glow: Generative Flow with Invertible 1x1 Convolu-](#)
1052 [tions](#). In *Proceedings of NeurIPS*, 2018. [0/2]
- 1053 • Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, David K. Duvenaud. [Neural Ordinary](#)
1054 [Differential Equations](#). In *Proceedings of NeurIPS*, 2018. [0/1]
- 1055 • Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, Jure Leskovec.
1056 [Hierarchical Graph Representation Learning with Differentiable Pooling](#). In *Proceedings of*
1057 *NeurIPS*, 2018. [0/1]
- 1058 • Ricky T. Q. Chen, Xuechen Li, Roger B. Grosse, David K. Duvenaud. [Isolating Sources of](#)
1059 [Disentanglement in Variational Autoencoders](#). In *Proceedings of NeurIPS*, 2018. [0/1]
- 1060 • Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, Baoquan Chen. [PointCNN:](#)
1061 [Convolution On X-Transformed Points](#). In *Proceedings of NeurIPS*, 2018. [0/1]
- 1062 • Arthur Jacot, Franck Gabriel, Clement Hongler. [Neural Tangent Kernel: Convergence and](#)
1063 [Generalization in Neural Networks](#). In *Proceedings of NeurIPS*, 2018. [0/1]
- 1064 • Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, Bryan
1065 Catanzaro. [Video-to-Video Synthesis](#). In *Proceedings of NeurIPS*, 2018. [0/1]

- 1066 • Yuanzhi Li, Yingyu Liang. [Learning Overparameterized Neural Networks via Stochastic](#)
1067 [Gradient Descent on Structured Data](#). In *Proceedings of NeurIPS*, 2018. [0/1]
- 1068 • Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, Aleksander Madry.
1069 [Adversarially Robust Generalization Requires More Data](#). In *Proceedings of NeurIPS*, 2018.
1070 [0/2]
- 1071 • Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, Aleksander Madry. [How Does Batch](#)
1072 [Normalization Help Optimization?](#) In *Proceedings of NeurIPS*, 2018. [0/1]
- 1073 • Harini Kannan, Alexey Kurakin, Ian Goodfellow. [Adversarial Logit Pairing](#). In *Proceedings*
1074 *of NeurIPS**, 2018. [0/2]
- 1075 • Ofir Nachum, Shixiang (Shane) Gu, Honglak Lee, Sergey Levine. [Data-Efficient Hierarchi-](#)
1076 [cal Reinforcement Learning](#). In *Proceedings of NeurIPS*, 2018. [0/3]
- 1077 • Prateek Jain, Raghu Meka, Inderjit Dhillon. [Guaranteed Rank Minimization via Singular](#)
1078 [Value Projection](#). In *Proceedings of NeurIPS*, 2010. [0/1]
- 1079 • Hanna Wallach, David Mimno, Andrew McCallum. [Rethinking LDA: Why Priors Matter](#).
1080 In *Proceedings of NeurIPS*, 2009. [0/4]
- 1081 • Geoffrey E. Hinton, Russ R. Salakhutdinov. [Replicated Softmax: an Undirected Topic](#)
1082 [Model](#). In *Proceedings of NeurIPS*, 2009. [0/1]
- 1083 • Daniel J. Hsu, Sham M. Kakade, John Langford, Tong Zhang. [Multi-Label Prediction via](#)
1084 [Compressed Sensing](#). In *Proceedings of NeurIPS*, 2009. [0/1]
- 1085 • Youngmin Cho, Lawrence Saul. [Kernel Methods for Deep Learning](#). In *Proceedings of*
1086 *NeurIPS*, 2009. [0/1]
- 1087 • Kurt Miller, Michael Jordan, Thomas Griffiths. [Nonparametric Latent Feature Models for](#)
1088 [Link Prediction](#). In *Proceedings of NeurIPS*, 2009. [0/3]
- 1089 • Ian Goodfellow, Honglak Lee, Quoc Le, Andrew Saxe, Andrew Ng. [Measuring Invariances](#)
1090 [in Deep Networks](#). In *Proceedings of NeurIPS*, 2009. [0/1]
- 1091 • Vinod Nair, Geoffrey E. Hinton. [3D Object Recognition with Deep Belief Nets](#). In *Proceed-*
1092 *ings of NeurIPS*, 2009. [0/1]
- 1093 • Martin Zinkevich, John Langford, Alex Smola. [Slow Learners are Fast](#). In *Proceedings of*
1094 *NeurIPS*, 2009. [0/1]
- 1095 • Ryan Mcdonald, Mehryar Mohri, Nathan Silberman, Dan Walker, Gideon Mann. [Efficient](#)
1096 [Large-Scale Distributed Training of Conditional Maximum Entropy Models](#). In *Proceedings*
1097 *of NeurIPS*, 2009. [0/1]
- 1098 • Corinna Cortes, Mehryar Mohri, Afshin Rostamizadeh. [Learning Non-Linear Combinations](#)
1099 [of Kernels](#). In *Proceedings of NeurIPS*, 2009. [0/1]
- 1100 • Laurent Jacob, Jean-philippe Vert, Francis Bach. [Clustered Multi-Task Learning: A Convex](#)
1101 [Formulation](#). In *Proceedings of NeurIPS*, 2008. [0/1]
- 1102 • Kamalika Chaudhuri, Claire Monteleoni. [Privacy-preserving logistic regression](#). In *Proceed-*
1103 *ings of NeurIPS*, 2008. [0/3]
- 1104 • Byron M. Yu, John P. Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V. Shenoy,
1105 Maneesh Sahani. [Gaussian-process factor analysis for low-dimensional single-trial analysis](#)
1106 [of neural population activity](#). In *Proceedings of NeurIPS*, 2008. [0/3]
- 1107 • Ilya Sutskever, Geoffrey E. Hinton, Graham W. Taylor. [The Recurrent Temporal Restricted](#)
1108 [Boltzmann Machine](#). In *Proceedings of NeurIPS*, 2008. [0/1]
- 1109 • Wenyuan Dai, Yuqiang Chen, Gui-rong Xue, Qiang Yang, Yong Yu. [Translated Learning:](#)
1110 [Transfer Learning across Different Feature Spaces](#). In *Proceedings of NeurIPS*, 2008. [0/3]
- 1111 • Yishay Mansour, Mehryar Mohri, Afshin Rostamizadeh. [Domain Adaptation with Multiple](#)
1112 [Sources](#). In *Proceedings of NeurIPS*, 2008. [0/1]
- 1113 • Sham M. Kakade, Karthik Sridharan, Ambuj Tewari. [On the Complexity of Linear Predic-](#)
1114 [tion: Risk Bounds, Margin Bounds, and Regularization](#). In *Proceedings of NeurIPS*, 2008.
1115 [0/1]

- 1116 • Francis Bach. Exploring Large Feature Spaces with Hierarchical Multiple Kernel Learning.
1117 In *Proceedings of NeurIPS*, 2008. [0/1]
- 1118 • Ijaz Akhter, Yaser Sheikh, Sohaib Khan, Takeo Kanade. Nonrigid Structure from Motion in
1119 Trajectory Space. In *Proceedings of NeurIPS*, 2008. [0/1]
- 1120 • Prateek Jain, Brian Kulis, Inderjit Dhillon, Kristen Grauman. Online Metric Learning and
1121 Fast Similarity Search. In *Proceedings of NeurIPS*, 2008. [0/1]
- 1122 • Duy Nguyen-tuong, Jan Peters, Matthias Seeger. Local Gaussian Process Regression for
1123 Real Time Online Model Learning. In *Proceedings of NeurIPS*, 2008. [0/1]
- 1124 • Lester Mackey. Deflation Methods for Sparse PCA. In *Proceedings of NeurIPS*, 2008. [0/1]