9 de abril de 2020

Minicorpus de Abstracts

Pedro Virgilio Basilio Jeronymo pedro.jeronymo@usp.br
Nº USP 8910559

1. Minicorpus

1.1 Setting

Argue about the topics prominence

> **The theory of** reinforcement learning **provides a normative account**, **deeply rooted in** psychological and neuroscientific perspectives on animal behaviour, of how agents may optimize their control of an environment.
>
> **Tags:** *introduce*

> Recurrent neural networks (RNNs) **are a powerful model for** sequential data.
>
> **Tags:** *describe*

> Topic models, such as latent Dirichlet allocation (LDA), **can be useful tools for the statistical analysis of** document collections and other discrete data.
>
> **Tags:** *describe*

> Recurrent neural networks (RNNs) **stand at the forefront of many recent developments in** deep learning.
>
> **Tags:** *compare*

> In many multivariate domains, **we are interested in** analyzing the dependency structure of the underlying distribution, e.g., whether two variables are in direct interaction.
>
> **Tags:** *introduce*

> **Estimating** influence on social media networks **is an important practical and theoretical problem** (...)
>
> **Tags:** *introduce*

Inductive transfer learning **has greatly impacted** computer vision (...)

Tags: *introduce*

Adversarial examples **pose** security **concerns because** they could be used to perform an attack on machine learning systems, even if the adversary has no access to the underlying model.

Tags: *introduce*

Familiarize terms, objects, or processes

End-to-end training **methods such as** Connectionist Temporal Classification **make it possible to** train RNNs for sequence labelling problems where the input-output alignment is unknown.

Tags: *describe*

The LDA **model assumes that** the words of each document arise from a mixture of topics, **each of which is** a distribution over the vocabulary.

Tags: *describe*

We can represent dependency structures using Bayesian network models. To analyze a given data set, Bayesian model selection **attempts to find** the most likely (MAP) model, and uses its structure to answer these questions.

Tags: *describe*

Skin cancer, the most common human malignancy, is primarily diagnosed visually, **beginning with an** initial clinical screening **and followed potentially by** dermoscopic analysis, a biopsy and histopathological examination.

Tags: *enumerate*

An adversarial example is a sample of input data which has been modified very slightly in a way **that is intended to** cause a machine learning classifier to misclassify it. **In many cases**, these modifications can be so subtle that a human observer does not even notice the modification at all, **yet the** classifier **still makes a mistake**.

Tags: *contrast*

Cite previous research results

Remarkably, humans and other animals **seem to solve this problem through a harmonious combination of** reinforcement learning and hierarchical sensory processing systems, **the former evidenced by a wealth of** neural data **revealing notable parallels between the** phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms.

Tags: *describe, compare, confirm*

The combination of these methods with the Long Short-term Memory RNN architecture **has proved particularly fruitful**, **delivering state-of-the-art results in** cursive handwriting recognition.

Tags: *describe, confirm*

Recent results at the intersection of Bayesian modelling and deep learning **offer a** Bayesian **interpretation of common** deep learning **techniques such as** dropout.

Tags: *describe*

Whereas before 2006 **it appears that** deep multilayer neural networks **were not successfully** trained, **since then** several algorithms **have been shown to successfully** train them, **with experimental results showing the superiority of** deeper **vs** less deep architectures. **All these experimental results were obtained with** new initialization or training mechanisms.

Tags: *contrast, confirm, compare*

Introduce hypotheses

This grounding of dropout **in** approximate Bayesian inference **suggests an extension of the theoretical results, offering insights into the use of** dropout **with** RNN **models**.

Tags: *introduce*

Deep convolutional neural networks (CNNs) **show potential for** general and highly variable tasks across many fine-grained object categories.

Tags: *confirm*

1.2 Gap

Cite problems/difficulties

However RNN **performance in** speech recognition **has so far been disappointing, with better results returned by** deep feedforward networks.

Tags: *compare*

A limitation of LDA **is the inability to model** topic correlation **even though, for example,** a document about genetics is more likely to also be about disease than X-ray astronomy. **This limitation stems from the use of** the Dirichlet distribution to model the variability among the topic proportions.

Tags: *contrast, compare*

Yet a major difficulty with these models is their tendency to overfit, with dropout shown to fail when applied to recurrent layers.

Tags: *contrast*

However, when the amount of available data is modest, **there might be** many models that have non-negligible posterior.

Tags: *contrast*

...**especially because this** new medium **is widely exploited as a platform for** disinformation and propaganda.

Tags: *describe*

Automated classification of skin lesions using images **is a challenging task owing to the** fine-grained variability in the appearance of skin lesions.

Tags: *describe*

...**but existing approaches in** NLP **still require** task-specific modifications and training from scratch.

Tags: *compare*

Most existing machine learning classifiers are highly vulnerable to adversarial examples.

Tags: *compare*

Cite needs/requirements

To use reinforcement learning **successfully in situations approaching** real-world complexity, **however**, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations.

Tags: *introduce, contrast*

Thus, we want compute the Bayesian posterior of a feature, i.e., the total posterior probability of all models that contain it.

Tags: *introduce*

Cite missing issues in previous research

While reinforcement learning agents **have achieved some successes in a variety of domains**, **their applicability has previously been limited to domains in which** useful features can be handcrafted, **or to domains with** fully observed, low-dimensional state spaces.

Tags: *compare*

Up to now, all previous work have assumed a threat **model in which the** adversary can feed data directly into the machine learning classifier. **This is not always the case for** systems operating in the physical world, **for example those which are using** signals from cameras and other sensors as an input.

Tags: *contrast, compare*

1.3 Purpose

Indicate main purpose

In this paper, we propose a new approach for this task.

Tags: *define*

Introduce purpose with methods

Here we use recent advances in training deep neural networks **to develop a novel** artificial agent, **termed a** deep Q-network, **that can** learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning.

Tags: *describe*

In this paper we develop the correlated topic model (CTM), **where the** topic proportions exhibit correlation via the logistic normal distribution [J. Roy. Statist. Soc. Ser. B 44 (1982) 139–177].

Tags: *describe*

This paper investigates deep recurrent neural networks, **which combine the** multiple levels of representation **that have proved so effective in** deep networks with the flexible use of long range context that empowers RNNs.

Tags: *describe*

We apply this new variational inference based dropout technique in LSTM and GRU models, **assessing it on** language modelling and sentiment analysis tasks.

Tags: *describe, introduce*

This paper introduces a novel approach to influence estimation on social media networks **and applies it to the** real-world problem of characterizing active influence operations on Twitter during the 2017 French presidential elections.

Tags: *describe, introduce*

Here we demonstrate classification of skin lesions using a single CNN, trained end-to-end from images directly, using only pixels and disease labels as inputs.

Tags: *describe*

We propose Universal Language Model Fine-tuning (ULMFiT), an effective transfer learning **method that can be applied to** any task in NLP, **and introduce techniques that are key for** fine-tuning a language model.

Tags: *introduce, describe*

Introduce purpose with results

This paper shows that even in such physical world **scenarios**, machine learning systems are vulnerable to adversarial examples.

Tags: *introduce*

Introduce purpose with Gap

Our objective here is to understand better why standard gradient descent from random initialization is doing so poorly with deep neural networks, **to better understand these recent relative successes and help design better** algorithms **in the future**.

Tags: *introduce*

1.4 Methods and Materials

Describe methods and materials

We tested this agent **on the challenging domain of** classic Atari 2600 games.

Tags: *describe*

We apply the CTM to the articles from Science published from 1990–1999, a data set that comprises 57M words.

Tags: *describe*

We then use this result as the basis for an algorithm that approximates the Bayesian posterior of a feature. **Our approach uses a** Markov Chain Monte Carlo (MCMC) method, but over orders rather than over network structures.

Tags: *describe*

The new influence estimation **approach attributes impact by accounting for** narrative propagation over the network. **This grounding of** network causal inference framework **applied to** data **arising from** graph sampling and filtering.

Tags: *introduce, describe*

We train a CNN using a dataset of 129,450 clinical images—two orders of magnitude larger than previous datasets — consisting of 2,032 different diseases. **We test its performance against** 21 board-certified dermatologists on biopsy-proven clinical images with two critical binary classification use cases: keratinocyte carcinomas versus benign seborrheic keratoses; and malignant melanomas versus benign nevi. The first case represents the identification of the most common cancers, the second represents the identification of the deadliest skin cancer.

Tags: *describe*

We demonstrate this by feeding adversarial images obtained from cell-phone camera to an ImageNet Inception classifier and measuring the classification accuracy of the system.

Tags: *describe*

Introduces new method contrasting with existing approaches

This causal **framework** infers the difference in outcome as a function of exposure, **in contrast to existing approaches that** attribute impact to activity volume or topological features, **which** do not explicitly measure nor necessarily indicate actual network influence.

Tags: *contrast*

Cramér-Rao estimation bounds **are derived for** parameter estimation as a step in the causal analysis, **and used to achieve** geometrical **insight on the** causal inference **problem**.

Tags: *describe*

1.5 Main Results

Describe the results

> **We demonstrate that the** deep Q-network agent, receiving only the pixels and the game score as inputs, **was able to surpass the performance of all previous** algorithms **and achieve a level comparable to that of a** professional human games tester across a set of 49 games, using the same algorithm, network architecture and hyperparameters.

Tags: *describe, compare*

> **When** trained end-to-end with suitable regularisation, **we find that** deep Long Short-term Memory RNNs **achieve a** test set error of 17.7% **on the** TIMIT phoneme recognition **benchmark** (...)

Tags: *describe*

> **We derive a** fast variational inference algorithm for approximate posterior inference in this model (...)

Tags: *introduce*

> **The new approach outperforms existing techniques, and to the best of our knowledge improves on the** single model state-of-the-art in language modelling with the Penn Treebank (73.4 test perplexity)

Tags: *compare*

> **We present empirical results on** synthetic and real-life datasets **that compare our approach to** full model averaging (when possible), to MCMC over network structures, and to a non-Bayesian bootstrap approach.

Tags: *compare*

> **The ability to** infer high causal influence **is demonstrated on** real-world social media accounts that are later independently confirmed to be either directly affiliated or correlated with foreign influence operations using evidence supplied by the U.S. Congress and journalistic reports.

Tags: *describe*

> **Our method significantly outperforms the state-of-the-art on** six text classification tasks, reducing the error by 18-24% on the majority of datasets. **Furthermore, with only** 100 labeled examples, **it matches the performance of** training from scratch on 100x more data.

Tags: *compare*

Outline the results

We first show how to efficiently compute a sum over the exponential number of networks that are consistent with a fixed order over network variables. **This allows us to** compute, for a given order, both the marginal probability of the data and the posterior of a feature.

Tags: *enumerate, describe*

The CNN **achieves performance on** par with all tested experts across both tasks, **demonstrating an** artificial intelligence **capable of** classifying skin cancer **with a level of competence comparable to** dermatologists.

Tags: *describe, compare*

Comments about the results

...**which to our knowledge is the best** recorded score.

Tags: *compare*

...**which is complicated by the fact that** the logistic normal is not conjugate to the multinomial.

Tags: *describe*

Describe methods and results simultaneously

We first observe the influence of the non-linear activations functions. **We find that** the logistic sigmoid activation **is unsuited for** deep networks with random initialization **because of its** mean value, which can drive especially the top hidden layer into saturation. **Surprisingly, we find that** saturated units can move out of saturation by themselves, **albeit** slowly, **and explaining the** plateaus sometimes seen when training neural networks. **We find that a new** non-linearity that saturates less **can often be beneficial**. **Finally, we study how** activations and gradients vary across layers and during training, **with the idea that** training may be more difficult when the singular values of the Jacobian associated with each layer are far from 1. B**ased on these considerations, we propose a new** initialization scheme **that brings** substantially faster convergence.

Tags: *enumerate, describe*

1.6 Conclusion

Outline conclusions

The CTM gives a better fit of the data than LDA, and **we demonstrate its use as an** exploratory tool of large document collections.

Tags: *describe*

Outfitted with deep neural networks, mobile devices **can potentially extend the reach of** dermatologists outside of the clinic. **It is projected that** 6.3 billion smartphone subscriptions will exist by the year 2021 (ref. 13) **and can therefore potentially** provide low-cost universal access to vital diagnostic care.

Tags: *describe*

We find that a large fraction of adversarial examples are classified incorrectly even when perceived through the camera.

Tags: *confirm*

Outline contributions/importance of research

This work bridges the divide between high-dimensional sensory inputs and actions, **resulting in the first** artificial agent **that is capable of** learning to excel at a diverse array of challenging tasks.

Tags: *confirm*

This extends our arsenal of variational **tools in** deep learning.

Tags: *describe*