

Pushing the technical frontier: From overwhelmingly large data sets to machine learning

Viviana Acquaviva

Physics Department, New York City College of Technology, 300 Jay Street,
Brooklyn NY 11201

email: vacquaviva@citytech.cuny.edu

Abstract. This paper summarizes my thoughts, given in an invited review at the IAU symposium 341 “Challenges in Panchromatical Galaxy Modelling with Next Generation Facilities”, about how machine learning methods can help us solve some of the big data problems associated with current and upcoming large galaxy surveys.

Keywords. Surveys, Big Data, Machine Learning

1. It’s everywhere!

Machine learning and data mining methods have become ubiquitous in all fields of science, including Astronomy, in the last decade. Figure 1 shows the result of a search on the NASA-ADS archive service for papers that contain the words “Machine Learning” either in the title or in the abstract, in the last 12 years, in the Astronomy and Physics archives; the growth trend is obvious. But of course it’s not just papers; during this time, we have seen the emergence of dedicated conferences and sessions in conferences (such as this one), scientific associations, perhaps best described in the Astrostatistics and Astroinformatics Portal (ASAIP) [website](#), special initiatives (for example the [COsmo INitiative](#)), dedicated programming packages, such as AstroPy ([Astropy Collaboration 2013, 2018](#)) and AstroML ([Vanderplas et al. 2012](#)), books, ad-hoc review panels from the National Science Foundation, and, of course, jobs. The most recent ASAIP blog post (10/2018) estimates that 15–20% of faculty job ads on the AAS job register have a strong component in statistics, machine learning, or data science. But what is machine learning? And most importantly, what can we do with it?

2. What can we do with ML?

Machine learning is the art of teaching a machine to make informed decisions. Examples of its applications include recognizing and characterizing objects based on similarities and differences, picking up patterns, and distinguishing signal from noise. The following is a (certainly incomplete) list of important applications of machine learning methods to Astronomy.

Save time. This is perhaps the chief application of machine learning in a Universe as vast as ours. A typical example is the classification of galaxies according to their morphological type; [de la Calleja & Fuentes \(2004\)](#) is, to my knowledge, the first example, but there have been many more since then (e.g. [Katebi et al. \(2019\)](#) and references therein); figure 2 above shows an example from [Dieleman et al. \(2015\)](#). In this case, trained humans are the best classifiers; we have an uncanny ability to generalize just from a few examples,




Figure 1. Machine learning related papers on the NASA/ADS archive from 2006 to today.




Figure 2. Dieleman *et al.* (2015) tackled the problem of galaxy morphology classification using deep learning.

and to weed out unimportant features. However, when the data sets become large (millions of objects), there are only two solutions: deploy the power of citizen science, with successful programs such as Galaxy Zoo (Lintott *et al.* 2008), or use machine learning. In a supervised approach, one would ask expert classifiers to provide “solutions” (estimated galaxy types) for a subset of objects, the learning set, and then teach the machine to learn from those examples; in an unsupervised approach, the machine determines its own set of galaxy types directly from the data, then assigns each object to a class (Hocking *et al.* 2018).

Provide a data-driven alternative to simplistic models. The alternative to machine learning is usually a model-based approach in which one creates intuition-based models of the objects, typically dependent on some input variables that are the parameters of the model. For example, when modeling galaxies, we use stellar population models paired with parameters such as stellar mass, stellar assembly history, metallicity, dust abundance and properties, and so on. The models are obviously an extreme simplification of complex objects such as real galaxies (in my own cheesy words: galaxies are not spheres, see Fig. 3). As a result, when we build likelihood functions and we infer




Figure 3. Alas, galaxies are not spheres.

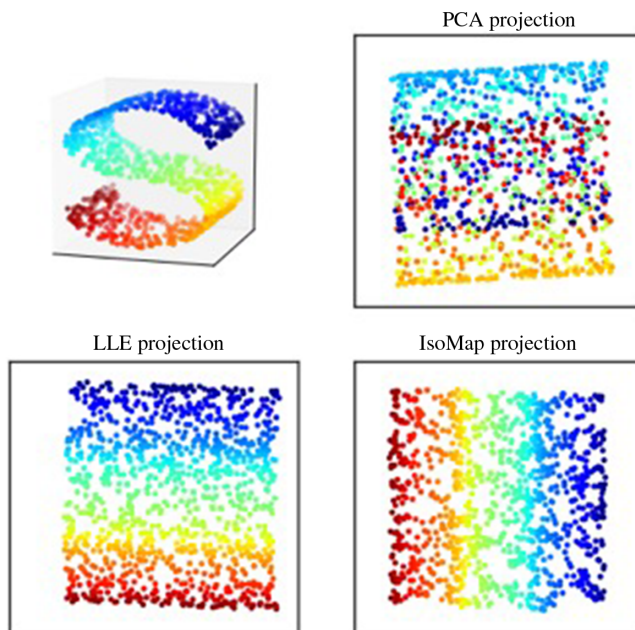


Figure 4. An example of $3D \rightarrow 2D$ dimensionality reduction, with various degrees of information loss. Figure from [the sklearn website](#).

probability distribution functions for our parameters, we bias our results by imposing our choice of parameters and priors. This model-related bias is something that machine learning can hope to circumvent, since the relation between input variables (observables) and output (measurements, or predictions) is purely data-driven.

Make problems more tractable. This is another field in which machine learning algorithms are quite amazing. One very simple but incredibly helpful application is dimensionality reduction, the process of identifying patterns in the data that allow one to reduce the size of the data space, making data easier to visualize, manipulate, and analyze, as shown in Fig. 4. Another great aspect is the increased resilience of ML methods to heterogeneous and sparse data. To stay within the regime of SED fitting, it is hard to create joint probability distributions to combine photometric and spectroscopic data, while for many ML methods, this problem just amounts to creating additional features.

Facilitate serendipitous discoveries. In 2017, a workshop called “Detecting the Unexpected” took place at the Space Telescope Science Institute. One of the most interesting talks (at least for me, following from afar) introduced me to the work of [Norris \(2017\)](#), who looked at the ten most important discoveries of the Hubble Space Telescope over the past decade. They pointed out that only a few key achievements, such as the use

Machine Learning	Model Fitting
<ul style="list-style-type: none"> • Data-driven (only as good as the data) • Usually generalizes poorly (model derived using some data can't be applied blindly to different data) • Interpretation is possible but might be non-trivial • Fast(er) • More robust/accommodating of mixed and missing data • Allows serendipitous discoveries 	<ul style="list-style-type: none"> • Intuition or model-driven (only as good as the scientist :)) • Generalizes well if physics is well understood • Easier to interpret • Might be computationally intensive • Dealing with heterogenous data often a pain in the neck • Leads to loss of information if models are too simplistic

Figure 5. Some advantages and disadvantages of the two approaches. My claim is that **synergy** is often the best strategy.

of Cepheids to improve the accuracy of the measurement of the Hubble constant, were planned in advance, while other fundamental advances in Astronomy, from the discovery of Dark Energy to the Hubble Deep Fields, were the result of serendipity. This - like many other examples - shows that one of the most important ways to prepare ourselves for further discoveries is to plan for what we don't know yet. In this respect, machine learning and data mining methods are really powerful; ML algorithms for pattern recognition and outlier detection are in general superior to model fitting, because the latter assumes, by construction, that we know what we are looking for.

3. Machine Learning versus Model Fitting: Who wore it better?

While I do, in many ways, “believe the hype” about machine learning, and I think it is important for young and old scientists alike to become proficient with these methods, I think it is important to realize that these techniques are just another set of tools in an astronomer/data scientist's skill set. I've often joked that machine learning methods are the “Numerical Recipes” of this decade. They enable us to do some science that would be difficult to do otherwise; but as always, it's part of a scientist's job to understand how to “attack” a difficult problem, with or without machine learning. As a general and incomplete rule, I think that when we understand the physics but we don't have data, we use model fitting, and when we don't understand the physics but we have data, we use machine learning; the figure 5 below summarizes some more of my thoughts on either approach, but the conclusion is that it's the synergy between these two approaches that is most powerful, because it gives us a strong cross-check of our problem-solving architecture.

4. Why now?

I hope that the discussion above provides a fair assessment of why machine learning methods are useful in Astronomy. But another question arises about timing: after all, none of these problems or methods are new, so why have we seen such a strong increase in ML-related science in the last few years? The key to that, in my opinion, is that we have a *super* data problem (my way of saying it's not just *big* data). Sure, data are becoming bigger; if we use etendue as zeroth-order tracer of the data volume generated by a survey, we can see in the left panel of Fig. 6 that the Large Synoptic Survey Telescope (LSST) dwarfs other surveys by a factor of 100 or more. However, we also have *better* data, because of the extended wavelength coverage, depth, and resolution of galaxy surveys. We have a *wide* data problem (to borrow a term I first heard from Alyssa Goodman at




Figure 6. An exciting and terrifying scenario: we have bigger data, better data, wide data, and new data. Figure sources, from left to right: LSST website, ESA website and Planck Collaboration (2016), and the Hubble website.

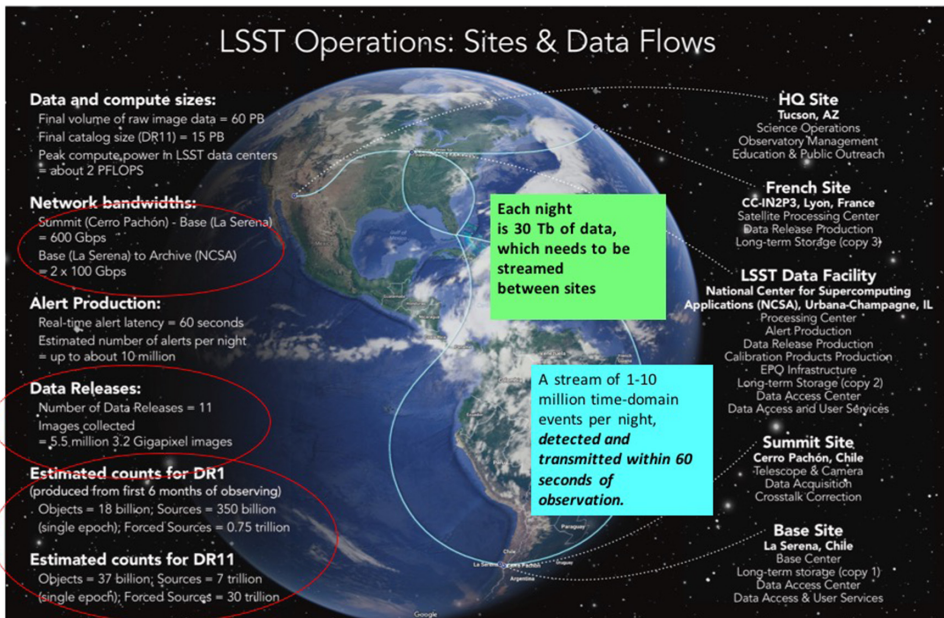


Figure 7. The “Big Data” problem of LSST. From the LSST website; annotations are mine.

a 2015 talk), where we have multiple representations of the same object, contributing information from different parts of the spectrum. And we have a *new* data problem, where we are bound to discover things that we didn’t think could exist, like hot Jupiters or super massive stars, or to unveil relationships that we don’t yet understand. It is in dealing with wide and new data that, in my opinion, machine learning methods really shine, and the timing is now because of upcoming observatories like LSST, the James Webb Space Telescope, and Euclid.

5. The example of LSST

The Large Synoptic Survey Telescope provides a great example of the type of big data problem that astronomers are facing, and it’s right around the corner, with a projected start date of 2023 (and an impressive record of being on schedule). Even the first of eleven data releases (DRs), anticipated to happen just six months after the beginning of science operations, will contain data for *eighteen billion* galaxies; DR11 will contain data for 37 billion galaxies, and will amount to the equivalent of 5.5 billion




Figure 8. Deep Learning algorithms (right) use multiple layers to characterize the relation between input and output. Figure from [this post](#).

Gigapixel-size images. That’s a lot of huge pictures! The problem of transferring and storing these data is only exacerbated by the fact that each night’s observations will amount to 30 Tb of data; within those, 1–10 million time-domain events will need to be streamed within 60 seconds, because real-time detection of unusual and serendipitous phenomena is essential; see Fig. 7. In one word: Ouch!

6. What’s next?

This section summarizes my thoughts about what are going to be the most pressing issues that we will face in the next five years or so, and what algorithms and techniques will, in my opinion, become most important in the “astronomy x data science” world. Of course, it’s a limited, biased view - if I were a ML algorithm, I would have low precision and low recall!

5.1. Too much information, too little space. One of the big issues we will face is going to be data storage. On the one hand, this applies to the raw data - for example, while searching and downloading the SDSS archive is a relatively manageable task for anyone with a laptop and a decent internet connection, the corpus of LSST data will certainly make this impossible. As a result, data compression, representation, and visualization techniques will become increasingly important. Examples of them range from the simple Principal Component Analysis to nonlinear techniques such as Self Organized Maps, t-distributed Stochastic Neighbor Embedding, Auto Encoders, and many others, some of which are described [here](#).

However, the tricky part is that different science applications will need different types of data, so the information loss associated to any technique may be perfectly acceptable for some projects, but not for others. The problem becomes even more pressing if one also considers *derived* products: catalogs, estimated parameters with their probability distributions, and so on. I found it telling that in my field of galaxy SED fitting, we spent a good fraction of the last decade arguing in favor of going beyond the best-fit model and using more sophisticated methods, such as various flavors of Markov Chain Monte Carlo methods, to explore the parameter spaces and generate full probability distributions. Instead, by the last CANDELS meeting (April 2018 at UC Riverside) the discussion had largely shifted towards what is the best way to make the products available, since full probability distributions are expensive to compute and store. I don’t really have a proposed solution, but I argue that we will all collectively need to discuss what products should be made available, and to develop/learn smart tools for data representation.

5.2. Deep Learning FTW. My prediction is that we will see a proliferation of deep learning methods (Fig. 8), a subset of machine learning methods in which the input/output relation is modeled through a series of intermediate steps (layers) that

Convolutional Neural Networks




Figure 9. A scheme of a convolutional neural network. The network includes convolution layers, in which complicated patterns are broken down into simple ones, creating data-driven features, pooling layers that are used for subsampling/dimensionality reduction, and a final fully connected layer that predicts the output. Figure courtesy of [Eugenio Colurciello](#).

make them really flexible and powerful in capturing features from raw data, moderating the need for feature selection and engineering, and modeling complex relationships.

The following are three types of neural networks that have already been shown to be very useful in Astronomy, and I expect that we'll hear a lot more about them in the next five years.

Convolutional Neural Networks. Convolutional Neural Networks (CNNs) are growing in popularity in many areas of Astronomy, typically as a means of analyzing 2D image data (e.g. [Dieleman et al. \(2015\)](#); [Tuccillo et al. \(2017\)](#); [Pettrillo et al. \(2017\)](#); [Fussell & Moews \(2018\)](#); [Huertas-Company et al. \(2018\)](#)). One powerful feature of CNNs is that they alternate different type of layers: convolution layers, used for feature extraction, and pooling layers, used for subsampling/dimensionality reduction. The final layer is a fully connected neural network, which generates the output. This structure, shown in Fig. 9 above, enables them to autonomously extract significant features from images, reducing the need for feature selection and feature engineering, and making them suitable to process astronomical images, in which information is often present in the form of gradients, breaks, and spatial correlations. A very nice, detailed discussion of how CNNs can master morphology classifications can be found in this [blog post](#).

Generative Adversarial Networks. Generative Adversarial Networks ([Goodfellow et al. 2014](#)) are built as a double system where one network attempts to build realistic examples of the objects, and the other one tries to detect whether they are fake or real. While the first applications were fraud detection, GANs have now been applied to a variety of problems, for example generating art in the style of a specific painter ([Zhu et al. 2017](#)); see Fig. 10 above. This [blog post](#) has a wonderful collection of examples. In Astronomy, GANs' recent claims to fame have been "upsampling" image resolution ([Ledig et al. 2016](#); [Schawinski et al. 2017](#)), and generating synthetic galaxy images that look realistic at the object and at the population level ([Fussell & Moews 2018](#)), as shown in Fig. 11. Both of these approaches have profound consequences for Astronomy, because they potentially lead to reduced observing times and to scaling down the processing time for large galaxy evolution simulations.

Recurrent Neural Networks. Recurrent neural networks, whose structure is shown in Fig. 12, have the characteristic of using not just the inputs, but the output of previous queries as part of their output generation. This characteristic makes them very suitable for all sorts of time series (time domain) problems, in which we care about not just the

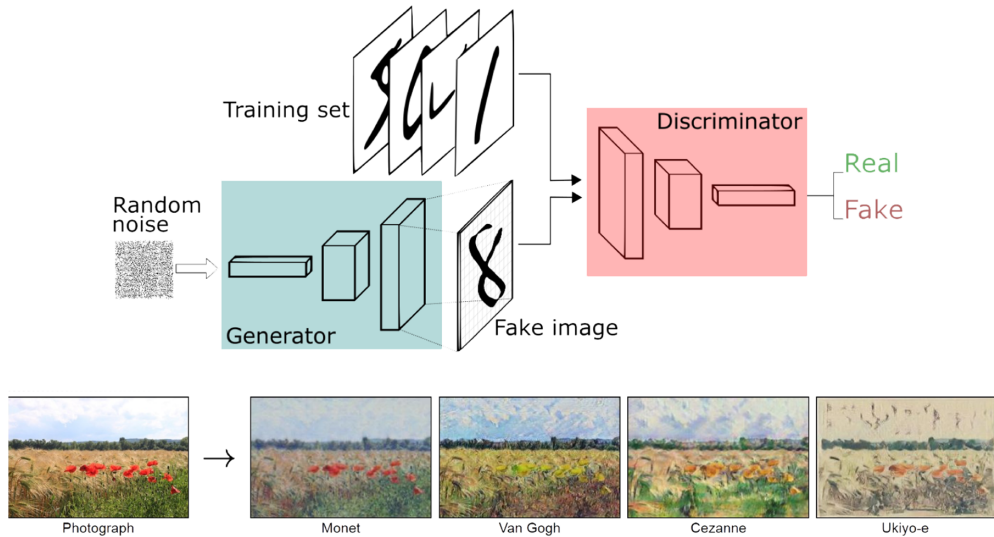


Figure 10. Top: Generative Adversarial Networks (GANs) learn how to generate realistic images thanks to their dual generator/discriminator structure. From [Skymind](#). Bottom: an example of style transfer, where the GAN processes the input photograph to generate art in the style of four different painters. From [Zhu et al. \(2017\)](#).

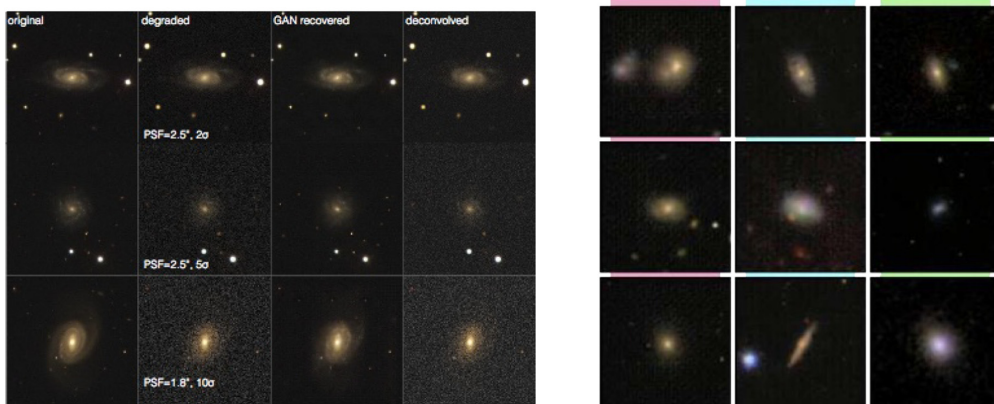


Figure 11. Left: An example of how GANs can be used to recover high-resolution galaxy images, from [Schawinski et al. \(2017\)](#). From left to right: the original SDSS image, the degraded image with a worse PSF and higher noise level (indicated), the image as recovered by the GAN, and for comparison, the result of a deconvolution. Right: Generating simulated galaxy images with GANs; the first two columns show synthetic objects, the last one real SDSS images. From v1 of [Fussell & Moews \(2018\)](#); further improvements in removing artifact features were solved by these authors through the development of a two-stage GAN approach.

characteristic of a single input instance but about a *sequence* of them. This is very helpful, among other cases, in outlier detection and hardware failure detection ([Naul et al. 2018](#); [Zhang & Zou 2018](#)).

5.3. (Sorry,) We need to talk to non-astronomers. An important consequence of the widespread impact of machine learning techniques in Astronomy is the fact that we need to develop interdisciplinary collaborations - for real! While this is in general always a good idea, I finally see it happening and it's in part due to the fact that data




Figure 12. Left: In a Recurrent Neural Network, the information flow is not just forward-fed, but the unit processors of the network can use information about the output of the network in previous steps. Figure by [Bruno Goncalves](#). Right: Because of this features, RNNs are good at modeling relationships between series of inputs and series of outputs. Source: this [blog post](#).

science seems to provide a common language for astronomers, statisticians, mathematicians, and computer scientists. One positive example is the fact that there have been several Astronomy/Physics challenges on the [Kaggle website](#), from the 2014 Galaxy Zoo challenge and Higgs boson challenge to recent ones like the [PLASTICC challenge](#) for light curves classification in LSST data and the ongoing [earthquake prediction challenge](#) promoted by Los Alamos National Lab. Personally, I think this is a wonderful thing.

5.4. What the heck are we doing? Model interpretability/Explainability. One drawback of powerful deep learning methods (but also of simpler models with heavily engineered features) is that they are difficult to interpret, hindering the process of producing generalizable insights. As a result, the last few years have seen the advent of interpretable machine learning (e.g. [Molnar 2018](#)), a discipline that seeks to increase the transparency of the machine’s decision-making process. Useful methodology tools for this task are LIME (Locally Interpretable Model-Agnostic Explanations, [Tulio Ribeiro et al. 2016](#)), which uses surrogate simpler models to map the input-output relationships of an “opaque” algorithm such as a CNN; SHAP (SHapley Additive exPlanations, [Lundberg et al. 2018](#)), which uses game theory to understand the interactions between different features (players) and their roles in the final decision, and Recourse Analysis [Ustun et al. 2018](#), which aims to build interpretability by understanding what feature modifications would change the outcome of a learning algorithm. The topic is controversial in the computer science community, and I believe that the trade-off between predictive power and interpretability should be decided on a discipline-by-discipline, or even case-by-case basis; see [Lipton \(2016\)](#) for a great discussion. Nonetheless, I imagine that many of these will be helpful in our quest to further our understanding of the physics of galaxies.

7. Conclusions

Machine learning and data mining methods provide an exciting opportunity to deal with our “super data” problems. Here are some ideas on how to leverage them:

- **Teach/learn this stuff early and often.** There is no doubt that any Astronomy graduate student can benefit from learning about data science methods, and there are infinite possibilities to do so. Many resources, such as classes and books, are available for free online; there are also dedicated initiatives such as the [LSST data science fellowship](#).
- **(Encourage mentees to) Code in Python.** This is essential, I think, for two reasons: one is convenience, because so many of the dedicated packages and support networks are available for Python, and second, because learning Python (and data science tools) will keep options open in case a career change is desired or necessary. I should add that some database experience and distributed computing experience is also quite desirable for this purpose.
- **(Encourage mentees to) Talk to people and go to conferences outside your field.** This year, the most inspiring talks I have heard were given in non-academic

settings. I want to acknowledge these people here: Mikio Braun (Zalando) and Adil Aijaz (Split Software) at Strata Data NYC, Manojit Nandi (Rocketrip) at PyGotham, Stephanie Yang (Foursquare) at a NYC Women in Machine Learning and Data Science meetup, and Francesca Lazzeri (Microsoft) at DataX NYC. This is of course driven by my own interests, but I think it's generally true that hearing about different topics and gathering fresh perspectives is exhilarating, and I have found these environments generally welcoming (not to mention the networking value).

- **(Encourage mentees to) Publish your code and make your results reproducible.** As the use of computational techniques becomes increasingly essential and algorithms become possibly more opaque, sharing code, even if it's not pretty, is the main way of keeping each other accountable and learning from each other. Plus, having an online presence on github is a good idea in view of career advancement, both in academia and in industry.

- **Rest assured that science is science.** Machine learning is not a magic wand; it's just a way of enhancing the tool set that we have. In fact, I think that it is even more essential to learn how to formalize, attack, solve, and test our solutions to problems now that we have more ways to solve them. In other words, the scientific method stands high and strong, and intuition and rigor continue, even more than before, to be the pillars of advancement in science.

References

- Astropy Collaboration. 2013, *Astron. Astrophys.*, 558, A33
- . 2018, *Astron. J.*, 156, 123
- de la Calleja, J. & Fuentes, O. 2004, *Mon. Not. R. Astron. Soc.*, 349, 87
- Dieleman, S., Willett, K. W., & Dambre, J. 2015, *Mon. Not. R. Astron. Soc.*, 450, 1441
- Fussell, L. & Moews, B. 2018, ArXiv e-prints, [1811.03081](https://arxiv.org/abs/1811.03081)
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. 2014, ArXiv e-prints, [1406.2661](https://arxiv.org/abs/1406.2661)
- Hocking, A., Geach, J. E., Sun, Y., & Davey, N. 2018, *Mon. Not. R. Astron. Soc.*, 473, 1108
- Huertas-Company, M., Primack, J. R., Dekel, A., Koo, D. C., Lapiner, S., Ceverino, D., Simons, R. C., Snyder, G. F., Bernardi, M., Chen, Z., Domínguez-Sánchez, H., Lee, C. T., Margalef-Bentabol, B., & Tuccillo, D. 2018, *Astrophys. J.*, 858, 114
- Katebi, R., Zhou, Y., Chornock, R., & Bunesco, R. 2019, *Mon. Not. R. Astron. Soc.*, 486, 1539
- Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., & Shi, W. 2016, arXiv e-prints, [arXiv:1609.04802](https://arxiv.org/abs/1609.04802)
- Lintott, C. J., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., Raddick, M. J., Nichol, R. C., Szalay, A., Andreescu, D., Murray, P., & Vandenberg, J. 2008, *Mon. Not. R. Astron. Soc.*, 389, 1179
- Lipton, Z. C. 2016, arXiv e-prints, [arXiv:1606.03490](https://arxiv.org/abs/1606.03490)
- Lundberg, S. M., Erion, G. G., & Lee, S.-I. 2018, ArXiv e-prints
- Molnar, C. 2018, Interpretable Machine Learning (<https://christophm.github.io/interpretable-ml-book/>), <https://christophm.github.io/interpretable-ml-book/>
- Naul, B., Bloom, J. S., Pérez, F., & van der Walt, S. 2018, *Nature Astronomy*, 2, 151
- Norris, R. P. 2017, *Publications of the Astron. Soc. of Australia*, 34, e007
- Petrillo, C. E., Tortora, C., Chatterjee, S., Vernardos, G., Koopmans, L. V. E., Verdoes Kleijn, G., Napolitano, N. R., Covone, G., Schneider, P., Grado, A., & McFarland, J. 2017, ArXiv e-prints, 1702, [arXiv:1702.07675](https://arxiv.org/abs/1702.07675)
- Planck Collaboration. 2016, *A&A*, 594, A10
- Schawinski, K., Zhang, C., Zhang, H., Fowler, L., & Santhanam, G. K. 2017, *Mon. Not. R. Astron. Soc.*, 467, L110
- Tuccillo, D., Huertas-Company, M., Decencièrre, E., & Velasco-Forero, S. 2017, in *IAU Symposium*, Vol. 325, *Astroinformatics*, ed. M. Brescia, S. G. Djorgovski, E. D. Feigelson, G. Longo, & S. Cavuoti, 191–196

- Tulio Ribeiro, M., Singh, S., & Guestrin, C. 2016, ArXiv e-prints, [1602.04938](#)
- Ustun, B., Spangher, A., & Liu, Y. 2018, ArXiv e-prints, [1809.06514](#)
- Vanderplas, J., Connolly, A., Ivezić, Ž., & Gray, A. 2012, in Conference on Intelligent Data Understanding (CIDU), 47 –54
- Zhang, R. & Zou, Q. 2018, Journal of Physics: Conference Series, 1061, 012012
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. 2017, arXiv e-prints, [arXiv:1703.10593](#)

Discussion

DAVID ROSARIO: With deep learning, there are examples of “emergent behaviour” where patterns and relationships are generated that were not represented in the training data. How could one identify such classes and either guard against or learn from them.

VIVIANA ACQUAVIVA: This can happen. Cross-validation and expert understanding of the data and the scientific problem are essential to understand emergent effects.