

Using Computational Creativity to Guide Data-Intensive Scientific Discovery

Kazjon Grace and Mary Lou Maher

{k.grace,m.maher}@uncc.edu

The University of North Carolina at Charlotte

Abstract

The generation of plausible hypotheses from observations is a creative process. Scientists looking to explain phenomena must invent hypothetical relationships between their dependent and independent variables and then design methods to verify or falsify them. Data-driven science is expanding both the role of artificial intelligence in this process and the scale of the observations from which hypotheses must be abduced. We adopt methods from the field of computational creativity – which seeks to model and understand creative behaviour – to the generation of scientific hypotheses. We argue that the generation of new insights from data is a creative process, and that a search for new hypotheses can be guided by evaluating those insights as creative artefacts. We present a framework for data-driven hypothesis discovery that is based on a computational model of creativity evaluation.

Introduction

Machine learning, for decades considered a standard tool throughout the sciences, is now being applied to discovery. Traditional approaches used AI to support hypotheses, building models that verify expected relationships between variables. AI systems are now regularly used to generate new and novel hypotheses from data (Colton and Steel 1999).

Another recent development in machine learning has been practical meta-optimisation: the use of optimisation approaches to tune the hyper-parameters of other optimisation problems. These sequential decision making methods – such as Bayesian optimisation (Snoek, Larochelle, and Adams 2012) – permit machine learning to be adopted by non-experts, essentially substituting required computational time for required expertise. Practical meta-optimisation (for which maturing tools are now freely available) facilitates machine learning as a hypothesis confirmation tool, simplifying the process of constructing effective predictive models of well-defined problems. This paper investigates the analogous problem: what would an automated tool look like for hypothesis discovery, rather than confirmation? How can methods like Bayesian Optimisation be extended to ex-

ploratory search, and what would be the objective function for such a tool?

This paper seeks answers to those questions in research on *creativity*, particularly in a domain of artificial intelligence called computational creativity: the computational modelling of behaviours that, if observed in humans, would be considered creative (Wiggins 2006). We argue that the exploration of possible hypotheses is a creative act, and that good hypotheses are creative artefacts. Building on research that defines a creative artefact as one that is *novel, valuable and surprising* (Maher and Fisher 2012; Grace et al. 2014a; 2014b) we define a framework for computational hypothesis discovery based on the search for hypotheses with the union of those three factors as the objective.

Computationally evaluating creativity

What motivates humans to deem artefacts creative? Given that we humans are responsible for constructing the notion of creativity, that question is analogous to defining creativity itself. A universally accepted definition is not readily available. Taylor (1988) gives dozens of attempts at defining creativity, from those that emphasise the creative process and aim to define it as a mental phenomenon to those that emphasise the social grounding of creative acts and aim to define it in terms of how society responds to creations. While these definitions were mostly founded with respect to artistic creativity, the same distinctions hold in science.

Newell, Shaw, and Simon (1959) define creative problem solving (an activity highly related to hypothesis generation) as occurring where one or more of the following conditions are met: the result has novelty and value for the solver or their culture, the thinking is unconventional, the thinking requires high motivation and persistence, and being ill-defined or requiring significant problem formulation.

Koestler (1964) and Boden (1990) offer related definitions based on the creative process rather than the resulting artefacts. Both are centred on the idea of a mental context for the domain in which the creative problem exists, a representation which Koestler calls a “matrix” and Boden a “conceptual space”. Koestler’s matrices are “patterns of ordered behaviour governed by fixed rules”, while Boden’s spaces are dimensioned by the “organising principles that unify and give structure to a given domain of thinking. Boden describes the pinnacle of creative thinking as the trans-

formation of the space of what is possible to include artefacts that were previously outside that space, while Koestler describes creativity as the blending, or “bisociation” of elements from two distinct mental frames of reference.

The transformative notions of the creative process postulated by Boden and Koestler are similar to those of unconventionality and reformulation raised by Newell. Each involves the relaxation of constraints and the adoption of elements – of process or product – considered outside the “norm”. The idea of transformation as a necessary component of creativity has attracted criticism for being difficult to operationalise (Wiggins 2006).

The notions of novelty and value first articulated by Newell are common among many of the definitions of creativity that can be found in the literature (Taylor 1988). This duality has recently been criticised by computational creativity researchers (Maher and Fisher 2012) as insufficient, with the argument that most computational models of novelty are based on solely on the difference between artefacts. Such objective measures do not capture the complexity and subtlety of *expectations* about artefacts and the way violating those expectations interacts with creativity.

To address this criticism we adopt the definition of a creative hypothesis as one that is novel, valuable and surprising, based on the framework described in (Maher and Fisher 2012). The additional surprise criterion incorporates the ideas of reformulation, transformation and bisociation raised in other definitions with the novelty/value duality of earlier definitions. Novelty (the degree to which an artefact is different from similar artefacts that already exist) and value (the degree to which an artefact is useful for its intended purpose) are near-universally accepted components of creativity. Surprise (the degree of unexpectedness of an artefact) is distinct from novelty as computational models of the latter capture only originality relative to the domain, and do not model notions of unexpectedness or the violation of trends. We adapt this three-factor measure of creativity for use in computational problem framing for analytics, where the “artefact” being evaluated for creativity is a relationship discovered within the data.

Creativity and the exploration of data

Creativity has long been identified as at the heart of the scientific endeavour (Taylor and Barron 1963). Hypothesis construction involves creatively transforming data in such a way that a new correlation emerges, and is as much a problem-framing activity as a problem-solving one. We propose that computational methods for evaluating the creativity of objects can be applied to hypothesis discovery, as good hypotheses are innovative, disruptive and require novel approaches to the problem – all hallmarks of creativity.

The use of AI methods for creative knowledge discovery dates back (to the authors’ knowledge) to AM and EURISKO (Lenat 1983), which generated new conjectures in mathematics, science and strategic wargames, among other disciplines. These programs possessed a model of the interestingness that guided their decision of what to explore next. One descendent of these programs, HR, expanded on interestingness, explicitly connecting interest-motivated discov-

ery to the notion of creativity (Colton and Steel 1999). Innovation analytics shares a goal with these early AI systems.

Innovation Analytics: A framework for computational hypothesis discovery

When the dataset becomes large and diverse, as is the case in any “big data” context, the problem of model selection rapidly becomes computationally prohibitive (Burnham and Anderson 2002). It is necessary to first select a subset of the overall data to relate to the desired dependent variables. Then it is often necessary to pre-process the selected variables in some way, by quantization, dimensionality reduction, or some other transformation. Only then can a supervised learning algorithm be applied to predict a relationship between the student attributes and the success attribute(s). At each step there are a number of hyperparameters to tune dependent on the data and algorithms selected.

These steps combine to produce a very large space of potential models to explore, with each point in that space representing a potential relationship between student data and student success. We propose a computational method for exploration of that space, using methods derived from computational creativity evaluation. Automated model selection is a rich and on-going area of research (Bozdogan 1987; Burnham and Anderson 2002, etc), from which this project is differentiated by its goal of scientific hypothesis discovery and its accompanying objective function of creativity.

We outline the innovation analytics framework in four parts: the computational models of value, novelty, and surprise which make up the creativity evaluation function, and the search process by which we propose to maximise them. We refer to *independent variables*, which make up the available data from which we intend to build a model, and *dependent variables*, about which we intend to form hypotheses conditional on the independent variables.

Generating influence factors

Innovation analytics is a search-based approach to hypothesis discovery. The space of hypotheses that explain the dependent variables will be explored for candidates which are deemed creative. Hypotheses will be represented as tree structures composed of selections over data, transformations of data, and predictive models applied to data, with each operation likely to incorporate several hyperparameters. We are currently investigating approaches for the search algorithm, and are currently considering Genetic Programming (Koza, Bennett III, and Stiffelman 1999), Bayesian Optimisation (Snoek, Larochelle, and Adams 2012) and Particle Swarm Optimisation (Poli, Kennedy, and Blackwell 2007) among others.

Innovation analytics defines a non-stationary search, with the objective function defined by novelty, value and surprise changing with the discovery of each new hypothesis. This drift in the objective function will guide the search trajectory of the system towards new, creative hypotheses.

We define innovation analytics generally as a search for hypotheses about the dependent variables expressed in terms of the independent variables, where the objective function is

the union of novelty, value and surprise. We will investigate several solutions for combining the three factors of novelty, value and surprise into a scalar value.

Assessing the value of a hypothesis

A creative artefact's value is the degree to which it is recognised as useful, and of the three aspects of creativity it parallels traditional effectiveness measures most closely – it is the union of value with novelty and surprise that defines creativity. Two major ways of assessing value have been investigated in the literature: assessing recognition socially (Sarkar and Chakrabarti 2011; Grace et al. 2014b) or assessing utility directly (Maher and Fisher 2012). For this project we adopt the latter approach, constructing a measure of the utility of a hypothesis: the strength of the predicted correlation between the data and a dependent variable.

Our value measure is based on the idea that a useful influence factor explains variance in the data that existing influence factors do not. Some factors may be highly predictive for a small subset of the data, while others may be only slightly better than random chance but apply near-universally – i.e. value is multiobjective, incorporating both predictive power and effect size. A valuable hypothesis is one that not just accurately explains many data points, but one which explains data that few other hypotheses do. A valuable design is not assessed by summing over performance dimensions, because there is no objective weighting function by which to do so.

We operationalise this notion by positing a “performance space” where each dimension is the explained variance of a single entry in the test set. In this case each point's explained variance is a dimension of performance, and valuable hypotheses exist in unpopulated parts of that space. Hypotheses can be clustered in this space, and then a distance measure applied between the candidate hypothesis and the nearest cluster. This represents how different the correlation is from similar correlations, a measure of its uniqueness in the space of explained variance of the dependent variable.

Hypotheses with high value will offer either stronger explanations of academic success than current hypotheses or explanations for groups of students that current hypotheses do not account for. As more hypotheses are discovered, the clusters and distance measures will adapt, reflecting the system's changing knowledge of the space. Hypotheses that would previously have been of high value will reduce in value if many highly similar hypotheses are discovered.

Assessing the novelty of a hypothesis

A creative artefact's novelty is the degree to which it differs from other artefacts of the same kind. Each hypothesis to be evaluated for creativity is represented by both the transformation of the data necessary to construct it and the correlation(s) predicted after that transformation is applied. Applied to identifying influence factors in data, this means hypotheses that transform the data in ways unlike other existing hypotheses, and in doing so find different kinds of correlations within the data.

(Maher and Fisher 2012) evaluate novelty by clustering observed artefacts and calculating the distance between an

object and its nearest cluster. They distinguish this from value in that novelty is calculated in the artefact description space, while value is calculated in the artefact performance space. An innovation analytics hypothesis is represented by a tree structure composed of selections, transformations and predictions about data. Using an edit distance based similarity metric we can define a space of possible hypotheses and cluster the known hypotheses within that space.

We will adopt a hierarchical clustering approach novelty evaluation, which is advantageous as the domain of possible hypotheses – like any domain of creative artefacts – is expected to exhibit structure at multiple levels. A hierarchical clustering approach requires making less assumptions about the number and scale of the clusters comprising the domain than would single-level clustering.

Novelty will be evaluated as the weighted sum of the distances between the new hypothesis and each of the clusters it is placed in, with an exponentially decreasing weight dependent on the depth within the tree of each cluster. This reflects that being an outlier in a more general cluster indicates significantly more novelty than being an outlier in a more specific one.

Assessing the surprisingness of a hypothesis

(Grace et al. 2014a) describe two kinds of creativity-relevant surprise: artefacts that violate observed *relationships*, and artefacts that require conceptual *restructuring* to comprehend. Both are contingent on the violation of strongly-held expectations about future events, but differ in the kinds of expectation being violated. Relational surprise arises from violating expectations about the value some attribute of an artefact will take based on the value of other attributes. Structural surprise arises from the expectation that existing knowledge about how a domain is structured is complete and correct, and is violated when knowledge must be restructured to accommodate new observations.

Structural surprise can be assessed via a hypothesis's impact on the hierarchical clustering used for novelty evaluation. When a new hypothesis is evaluated it perturbs the structure of the online hierarchical clustering of observed hypotheses, affecting how others a clustered in addition to itself. The scale of this perturbation can be measured using a delta function, measuring the difference between the cluster hierarchy before and after the new hypothesis is observed.

Relational surprise models the interactions between the independent and dependent attributes. Each independent attribute will be initially assumed to be statistically independent of each dependent attribute, and these assumptions will be updated with each iteration of the system to reflect the degree of interaction predicted by all available hypotheses. The expected influence between variables in future iterations can then be modelled using regression. Sudden increases in the relationship between variables – those that violate the trend predicted by past observations – will elicit surprise.

In both of these models, surprising hypotheses about relationships between data are those that force the system to reconsider its understanding of the problem domain. Structural surprise occurs when expectations about the different kinds of hypothesis that exist are violated, causing the sys-

tem to restructure its clustering of what is observed. Relational surprise occurs when expectations about the relationships between attributes and success are violated, causing the system to re-evaluate their expected future interdependence. A hypothesis which violates expectations in either or both cases is considered surprising by the system.

Proposed domain: Learning analytics

Large-scale educational institutions, whether they are predominantly campus-based like many comprehensive state universities, or predominantly online like MOOCs, graduate far fewer students than they enrol. The US national average graduation-within-six-years rate for public four-year universities is 46%¹. Completion rates among MOOC courses are lower still, with completion rates of less than 10% being the norm. Those figures are far in excess of what can be explained by attrition of the incapable or unmotivated.

Educational organizations are increasingly in possession of overwhelming data about student activities, both in traditional campus-based and large-scale online learning environments. Most educational analytics research (Campbell, DeBlois, and Oblinger 2007; Ferguson 2012) has focused on predicting overall academic success from intuitively related sources, primarily other academic data. Innovation analytics permits an exploration of heterogeneous data that is not intuitively related to academic performance. Potential data sources include online activity, campus purchasing habits, swipe card logs and club/society activity that are initially non-intuitive. In this domain it is necessary not only to be able to predict student success, but to construct scrutable, verifiable hypotheses on which policy can be based.

Conclusion

We frame the act of discovering new hypotheses as a creative process, and propose a novel method of representing that act computationally based on that framing. The approach is inspired by theories and models developed for evaluating creative products, providing a theoretical basis for generating and evaluating creative hypotheses. Creativity is operationalised as the intersection of novelty, value and surprise: good designs are novel, unexpected and valuable to their users, while good influence factors are novel, unexpected and valuable to the research community. Novel influence factors promote progress, unexpected influence factors challenge established norms, and valuable influence factors explain the data. When the computational system identifies a correlation it evaluates as being creative, it communicates to its users the discovery, the interpretation of the data needed to discover it, and a justification for its discovery. This project posits a new approach to data-intensive analytics in scientific research that goes beyond a focus on finding strong correlations in large datasets. This project will develop a framework for discovering and communicating creative hypotheses about educational data.

¹Statistics from the National Center for Education Statistics (nces.ed.gov)

References

- Boden, M. A. 1990. *The Creative Mind: Myths and Mechanisms*. Routledge.
- Bozdogan, H. 1987. Model selection and Akaike's information criterion (AIC): The general theory and its analytical extensions. *Psychometrika* 52(3):345–370.
- Burnham, K. P., and Anderson, D. R. 2002. *Model Selection and Multi-model Inference: a Practical Information-theoretic Approach*. Springer.
- Campbell, J. P.; DeBlois, P. B.; and Oblinger, D. G. 2007. Academic analytics: A new tool for a new era. *Educause Review* 42(4):40.
- Colton, S., and Steel, G. 1999. Artificial intelligence and scientific creativity. *Artificial Intelligence and the Study of Behaviour Quarterly* 102.
- Ferguson, R. 2012. Learning analytics: drivers, developments and challenges. *International Journal of Technology Enhanced Learning* 4(5):304–317.
- Grace, K.; Maher, M. L.; Fisher, D.; and Brady, K. 2014a. Modeling expectation for evaluating surprise in design creativity. In *Design Computing and Cognition 2014 (to appear)*.
- Grace, K.; Maher, M. L.; Fisher, D.; and Brady, K. 2014b. A data-intensive approach to predicting creative designs based on novelty, value and surprise. *International Journal of Design, Creativity and Innovation (to appear)*.
- Koestler, A. 1964. *The Act of Creation*. Hutchinson London.
- Koza, J. R.; Bennett III, F. H.; and Stiffelman, O. 1999. *Genetic programming as a Darwinian invention machine*. Springer.
- Lenat, D. B. 1983. Eurisko: a program that learns new heuristics and domain concepts: the nature of heuristics iii: program design and results. *Artificial intelligence* 21(1):61–98.
- Maher, M. L., and Fisher, D. H. 2012. Using AI to evaluate creative designs. In *2nd International conference on design creativity (ICDC)*, 17–19.
- Newell, A.; Shaw, J.; and Simon, H. A. 1959. *The processes of creative thinking*. Rand Corporation.
- Poli, R.; Kennedy, J.; and Blackwell, T. 2007. Particle swarm optimization. *Swarm intelligence* 1(1):33–57.
- Sarkar, P., and Chakrabarti, A. 2011. Assessing design creativity. *Design Studies* 32(4):348–383.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical bayesian optimization of machine learning algorithms. In *NIPS*, 2960–2968.
- Taylor, C. W., and Barron, F. E. 1963. *Scientific creativity: Its recognition and development*. John Wiley.
- Taylor, C. W. 1988. Various approaches to the definition of creativity. In Sternberg, R. J., ed., *The Nature of Creativity*. Cambridge University Press. 99–124.
- Wiggins, G. A. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.