SIGEVOlution

newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation

Volume 11
Issue 1

in this issue

Evolution is the New Deep Learning

EvoSTAR2018

Welcome to the first 2018 newsletter. In their article “Evolution is the New Deep Learning”, Risto Miikkulainen and his collaborators highlight the power of Evolutionary Computation (EC), which they are argue can transform Artificial Intelligence as Deep Learning (DL) has done recently. While DL can model what we already know, EC can create new solutions that do not yet exist. So EC is the next step for DL, they argue, and follow on by showcasing several exciting EC articles and demos as an inspiration for future success stories. Regarding EC forthcoming events, we report on EvoSTAR 2018 taking place this April in the beautiful city of Parma, hosting interesting keynotes and research talks; and the forthcoming 14th International Workshop on Global Optimization (LeGO), which will take place from 18-21th September 2018, at Leiden University, The Netherlands, with a submission deadline in April 30th. As ever, please get in touch if you would like to contribute an article for a future issue or have ideas for the newsletter

Emma Hart and **Gabriela Ochoa**, Editors

14th International Workshop on Global Optimization (LeGO)

The workshop will take place from 18-21th September 2018, at Leiden University, The Netherlands. Website, PC, and Details: http://moda.liacs.nl/LeGO

It will feature a Special Track on Multiobjective Global Optimization, organized by **Iryna Yevseyeva**, De Montfort University, UK

The workshop will be organized by Leiden University and the Int’l Society of Global Optimization (ISoGO).

Local Organization Team:

Michael Emmerich (LIACS, Leiden Univ.),

André Deutz (LIACS, Leiden University),

Sander Hille (Mathematical Institute, Leiden Univ.),

Invited speakers:

Sergiy Butenko, Professor at Texas A&M University, USA

Kaisa Miettinen, Professor at the University of Jyväskylä, Finland

Panos M. Pardalos, Distinguished Professor at the University of Florida, USA

Yaroslav D. Sergeyev, Distinguished Professor University of Calabria, Italy

Antanas Žilinskas, Professor at Vilnius University, Lithuania

We welcome papers on global optimization methods and applications: Submission of 4 page, double column paper until April 30th, 2018 (general track), May 18th special track.

AIP - Web of Science indexed proceedings & Special Issue in Journal on Global Optimization.

Front cover image supplied by **Risto Miikkulainen**, Professor of Computer Science at the University of Texas at Austin and VP of Research at Sentient Technologies.

Evolution is the New Deep Learning

By **Risto Miikkulainen, Babak Hodjat, Xin Qiu, Jason Liang, Elliot Meyerson, Aditya Rawal, and Hormoz Shahrzad.**

Sentient Technologies and The University of Texas at Austin

Deep learning (DL) has transformed much of AI, and demonstrated how machine learning can make a difference in the real world. Its core technology is gradient descent, which has been used in neural networks since the 1980s. However, massive expansion of available training data and compute gave it a new instantiation that significantly increased its power.

Like DL, Evolutionary Computation (EC) was introduced decades ago, and it is currently experiencing a similar boost from the available big compute and big data. However, EC addresses a different but equally far-reaching problem. While DL is focused on modeling what we already know, EC is focused on creating solutions that do not yet exist. In that sense, it is the next step up from DL: Whereas DL makes it possible to recognize e.g. new instances of objects and speech within familiar categories, EC makes it possible to discover entirely new objects and behaviors---those that maximize a given objective. EC does it not by following a gradient (like most DL and reinforcement learning approaches), but by doing massive exploration: using a population of candidates to search the space of solutions in parallel, emphasizing novel and surprising solutions. Thus, EC makes a host of new applications of AI possible: designing more effective and economical physical devices and software interfaces; discovering more effective and efficient behaviors for robots and virtual agents; creating more effective and cheaper health interventions, growth recipes for agriculture, and mechanical and biological processes.

However, despite such opportunities, EC is not well known in the mainstream AI and Machine Learning (ML) communities. At Sentient, we recently launched a website aiming to reach outside of our own EC community, and to build a case for the argument above. We are showcasing five newpapers, but also illustrate EC through 11 animated demos and 3 interactive demos. The goal is to get the word out, i.e. to get people in AI thinking and talking about evolutionary computation. Given recent research at OpenAI, Uber, DeepMind, and Google, there is already significant momentum building in this area: EC is on the verge of becoming the next Deep Learning.

We would like to encourage other groups to engage in similar efforts. The goal of this article is thus not to repeat the contents of the website, but to motivate the development of future efforts to get the word out. In particular, it showcases the three interactive demos built for the site, hoping that they can serve as inspiration for similar examples in the future. As general context for these demos, the website supports the "EC is the next DL" argument from three perspectives:

(1) EC can be used to discover DL architectures that are more complex and powerful than those designed by hand. Neuroevolution techniques are shown to advance state of the art in three DL benchmarks;

(2) EC can be used to build commercial applications. Sentient Ascend is showcased as such an application in discovering web interfaces that convert casual users into customers better than hand-designed interfaces. Interestingly, efforts to develop such applications reveal interesting scientific problems, and solving them advances the field of EC itself; and

(3) EC can be used to solve very hard problems, including those that have a large search space, high dimensionality, and/or are highly deceptive. Discovering minimal sorting networks is used as an example such problem.

The details of the technology in these areas are described in nine research papers. The website then complements these papers with 11 demos that animate how evolution searches for solutions. Success of EC sometimes looks like magic, at odds with statistical notions of significance and confidence. The main goal is to make it clear how population-based search differs from gradient descent and reinforcement learning: that populations allow exploring wider areas of search space, making it possible to discover more creative solutions. In addition, the demos aim to make EC techniques such as neural architecture search, noisy evaluation, winner selection, multiobjective optimization, and novelty search concrete and understandable.

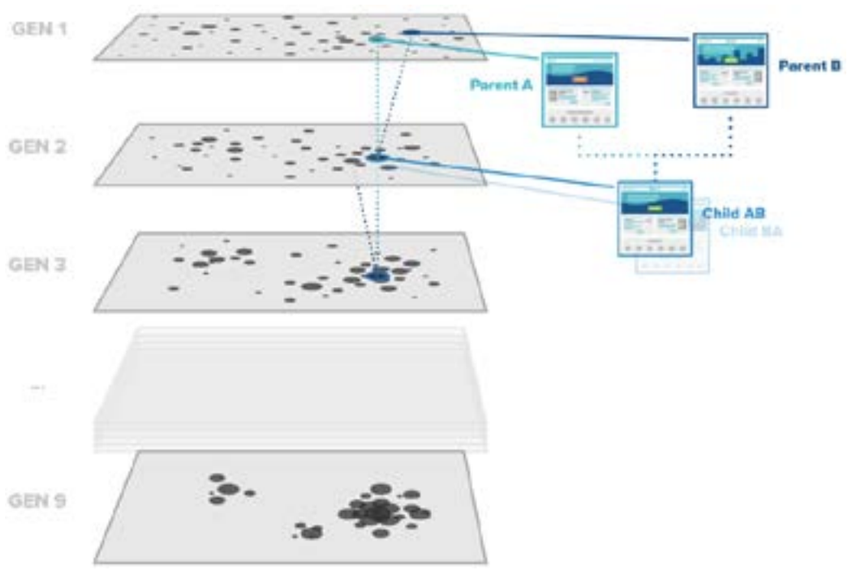

An important part of the argument is that EC can be used to improve the state of the art in several standard DL benchmarks: In that sense as well, it is the next step from DL. Beyond performance numbers, it turns out possible to illustrate each benchmark with an interactive demo that lets people understand what these discoveries mean. They are:

1. **Music Maker (Language Modeling)**

One of the benchmark tasks is Language Modeling, where the system is trained to predict the next word in a language corpus, i.e. a large collection of text such as several years of the Wall Street Journal. After the network has made its prediction, this input can be looped back into its input, and the network can generate an entire sequence of words. Interestingly, the same technique applies equally well to musical sequences, where it makes for a fun demo. The user inputs a few initial notes, and the system improvises an entire melody based on that starting point. Through neuroevolution, the design of the gated recurrent (Long Short-Term Memory or LSTM) nodes (i.e. the network's memory structure) was optimized to make the model more accurate in predicting the next note.

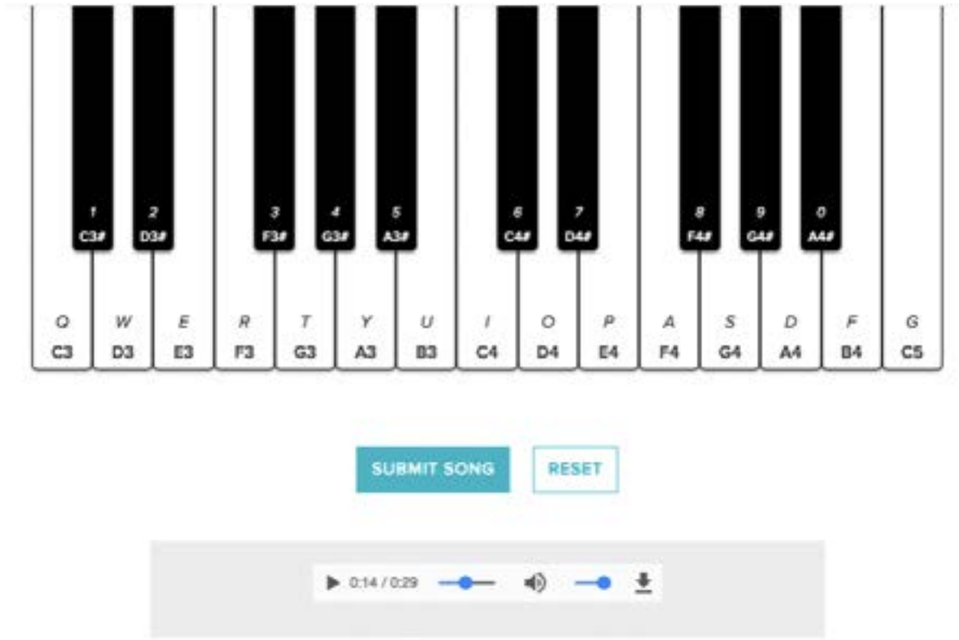

In the actual language modeling domain (i.e. predicting the next word in the Penn Tree Bank corpus), the benchmark is defined by perplexity points, a measurement of how well a probabilistic model can predict real samples. The lower the number the better, because we want the model to be less perplexed when predicting the next word in a sequence. In this case, the evolved structures beat the standard LSTM structure by 10.8 perplexity points. Remarkably, although several human-designed LSTM variations have been proposed, they did not improved performance much---before neural architecture search methods, LSTM structure was essentially unchanged for 25 years. Our neuroevolution experiments showed that it can, as a matter of fact, be improved significantly by adding more complexity, i.e. memory cells and more nonlinear, parallel pathways.

Why does this result matter? Language is a powerful and complex construct of human intelligence. Language modeling, i.e. predicting the next word in a text, is a benchmark that measures how well machine learning methods can learn language structure. It is therefore a surrogate for building natural language processing systems that includes speech and language interfaces, machine translation (such as Google Translate), and even medical data such as DNA sequences and heart-rate diagnosis. The better we can do in the language modelling benchmark, the better language processing systems we can build, using the same technology.

2. **Omni Draw (character recognition)**

Omniglot is a handwritten character recognition benchmark on recognizing characters in 50 different alphabets, including real languages like Cyrillic (written Russian), Japanese, and Hebrew, to artificial languages such as Tengwar (the written language in Lord of the Rings).

This demo showcases multitask learning, in which the model learns all languages at once and exploits the relationship between characters from different languages. So, for instance, the user inputs an image and the system outputs suggestions for different character matches in different languages, saying “this would be ‘X’ in Latin, ‘Y’ in Japanese, and ‘Z’ in Tengwar, etc.”---taking advantage of its understanding of the relationships between Japanese, Tengwar, and Latin to figure out which character is the best match. This approach differs from a single task learning environment where the model trains on one language at a time and cannot make the same connections across language data sets. In this Omniglot multitask character recognition problem, the neuroevolution approach reduced the error from 22% to 10%.

Omniglot is an example of a dataset that has relatively little data per language---for instance, it may have only a few characters in Greek but many in Japanese. It succeeds by using its knowledge of the relationships between languages to find solutions, hence, finding a solution in the face of missing or sparse data. Why is this important? For many real world applications, labeled data is expensive or dangerous to acquire (e.g., medical applications, agriculture, and robotic rescue), hence automatically designing models that exploit the relationships to similar or related data-

sets could, in a way, substitute the missing dataset and boost research capabilities. It is also an excellent demonstration of the power of neuroevolution: there are many ways in which the languages can be related, and evolution discovers the best ways to tie their learning together.

3. **Celeb Match (face understanding)**

The Celeb Match demo is also built upon multitask learning, but this time, with a large-scale data sets. The demo is based on the CelebA dataset, which consists of around 200,000 images of celebrities, each of which is labeled with 40 binary attributes such as “Male vs. Female”, “beard vs. no beard”, “glasses vs. no glasses”, etc. Each attribute induces a classification task that induces the system to detect and identify each attribute. As a fun add-on, we created a demo that turns this task around: The user can set the desired degree for each attribute, and the system finds the closest celebrity match, as determined by the evolved multitask learning network. For instance, if the current attribute settings result in an image of Brad Pitt, the user can increase “gray hair” to find which celebrity would be similar to Brad Pitt but with different hair.

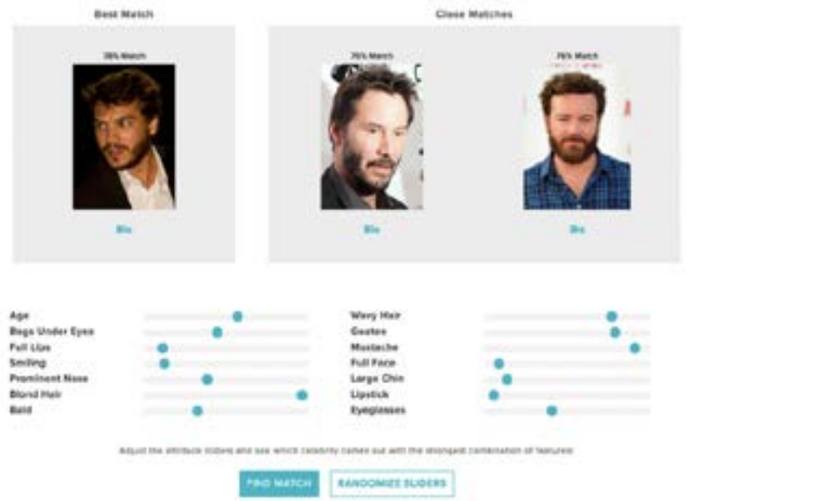

In this domain, the state-of-the-art benchmark is the test error across all attributes, i.e. whether the system detected the attribute correctly (male/female, young/mature, large eyes/small eyes), etc. In the CelebA multitask face classification domain, evolutionary computation was used to optimize the networks that detect these attributes, reducing error from 8.00% to 7.94% for an ensemble (an average of) three models. This technology is a step forward in the ability for AI to predict diverse attributes of people, places, and things in the physical world. Unlike networks trained to find similarities based on abstract, learned features, it makes the similarities semantic and interpretable.

Music Maker, Omni Draw, and Celeb Match, are just three examples of interactive demos that illustrate the power of EC. The idea of these demos is to entice people to learn more about the technology behind them on the website and papers. Through similar efforts, we hope that EC can be made concrete and familiar, and thus play the role it can play in building future AI systems.

EvoSTAR2018

EvoStar is the Leading European Event on Bio-Inspired Computation. This year it will take place in Parma, Italy from April 4th to 6th. EvoStar 2018 consists of four co-located conferences:

1. EuroGP 21st European Conference on Genetic Programming
2. EvoCOP 18th European Conference on Evolutionary Computation in Combinatorial Optimisation
3. EvoMUSART 7th International Conference (and 13th European event) on Evolutionary and Biologically Inspired Music, Sound, Art and Design
4. EvoApplications 21st European Conference on the Applications of Evolutionary and bio-inspired Computation including several tracks

The EvoSTAR 2018 program is now available at: http://www.evostar.org/2018/programme_overview.php

Invited Speakers

The event will host to wonderful invited speakers. The opening talk on Wednesday, April 4 is entitled Adversarial Dynamics: Understanding Them Now is Important as Ever, given by **Una-May O'Reilly** from AnyScale Learning for All Group, MIT Computer Science and Artificial Intelligence Lab, Cambridge, Massachusetts, US.

Abstract: The world abounds with adversarial relationships - opponents squaring off in contests and disputes. They are found in nature (predator vs prey) and in human society (gaming, wars, cyber security). They shape important outcomes and influence the character of our existence. I will describe my research agenda on the dynamics of adversarial relationships and understanding how optimization and evolutionary computation comprise a useful lens through which they can be examined.

Bio: Una-May O'Reilly leads the AnyScale Learning For All (ALFA) group at MIT. She has expertise in scalable machine learning, evolutionary algorithms, and frameworks for large-scale, automated knowledge mining, prediction and analytics. She educates the forthcoming generation of data scientists, teaching them how develop state of art techniques that address the challenges spanning data integration to knowledge extraction.

The closing talk on Friday, 6 April has the provocative title Evolution, Art and Sex, by **Penousal Machado** from the University of Coimbra, Coimbra, Portugal.

Abstract: Applying the techniques of Evolutionary Computation to artistic expression in both the visual arts and music over a period of several years has generated a timeline of what might be termed cultural artefacts. This suggested to me that the field - its past and current methodology, its trends and challenges - could be viewed through the prism of archeology. We will focus on one of the most intractable problems in the field: fitness assignment, analysing this challenge from the perspective of a human user interacting with an evolutionary system. How do Machine Learning, Evolutionary Computation and HCI techniques combine to create a Computer Aided Creativity system that responds to the artistic intentions of the user?

Bio: **Penousal Machado** leads the Cognitive and Media Systems group at the University of Coimbra. His research interests include Evolutionary Computation, Computational Creativity, and Evolutionary Machine Learning. In addition to the numerous scientific papers in these areas, his works have been presented in venues such as the National Museum of Contemporary Art (Portugal) and the "Talk to me" exhibition of the Museum of Modern Art, NY (MoMA).

EvoStar Outstanding Students 2018

This year, EvoStar wants to recognise the good work of students, who represent the future of our community. Papers that obtained an overall high review score and whose first author was a student registered to attend EvoStar, have been nominated for a Best Paper Award. Ten Outstanding Students were selected:

1. **Filipe Assunção**, CISUC, University of Coimbra, Coimbra, Portugal
2. **Mosab Bazargani**, Queen Mary University of London, London, United Kingdom
3. **Alexander Berman**, AI_am, Gothenburg, Sweden
4. **Sondre Engebraten**, FFI, Oslo, Norway
5. **Tiago Martins**, CISUC, University of Coimbra, Coimbra, Portugal
6. **Almuth Meier**, University of Oldenburg, Oldenburg, Germany
7. **Rafaela Priscila Cruz Moreira**, CEFET-MG, Belo Horizonte, Brazil
8. **Robert Jacob Smith**, Dalhousie University, Halifax, Canada
9. **Sara Tari**, University of Angers, Angers, France
10. **Sarah Louise Thomson**, University of Stirling, Stirling, United Kingdom

Best Paper Candidates

Each of the four co-located conferences will be giving a best-paper award. The nominees can be found here: http://www.evostar.org/2018/programme_bestpapers.php

HM 2019 - 11th International Workshop on Hybrid Metaheuristics

January 16-18, 2019 - Concepción, Chile

Website: http://hm2019.ing.udec.cl/
Submission: https://easychair.org/conferences/?conf=hm2019
Proceedings: Spinger LNCS

Metaheuristics are considered state-of-the-art methods for many hard optimization problems. At some point, however, it has become evident that the concentration on a sole metaheuristic is rather restrictive. A skilled combination of concepts from different optimization techniques can provide a more efficient behavior and a higher flexibility when dealing with real-world and large-scale problems. Hybrid Metaheuristics are techniques for optimization that combine different metaheuristics or integrate AI/OR techniques

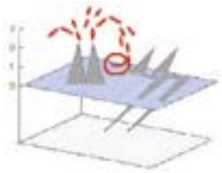

MESS 2018 - Metaheuristics Summer School

- from Design to Implementation -
21-25 July 2018, Taormina, Italy

**** APPLICATION DEADLINE: 15th April 2018 ****

MESS 2018 is aimed at qualified and strongly motivated MSc and PhD students; post-docs; young researchers, and both academic and industrial professionals to provide an overview on the several metaheuristics techniques, and an in-depth analysis of the state-of-the-art. As first edition, MESS 2018 wants to analyze all metaheuristics from its designing to its implementation. In particular, in MESS 2018 will be analyzed modern heuristic methods for search and optimization problems, as well as the classical exact optimization methods, seen also in the metaheuristics context.

** Short Talk and Poster Presentation

All participants may submit an abstract of their recent results, or works in progress, for presentation and having the opportunities for debate and interact with leaders in the field. Mini-Workshop Organizers and Scientific Committee will review the abstracts and will recommend for the format of the presentation (oral or poster). All abstracts will be published on the electronic hand-out book of the summer school.

The Abstracts must be submitted by *__April 15, 2018__*.

https://www.euro-online.org/websites/verolog/event/mess-2018-metaheuristics-summer-school/

First Workshop on Developmental Neural Networks

Overview

In nature, brains are built through a process of biological development in which many aspects of the network of neurons and connections change are shaped by external information received through sensory organs. Biological development mechanisms such as axon guidance and dendrite pruning have been shown to rely on neural activity. Despite this, most artificial neural network (ANN) models do not include developmental mechanisms and regard learning as the adjustment of connection weights, while some that do use development restrain it to a period before the ANN is used. It is worthwhile to understand the cognitive functions offered by development and to investigate the fundamental questions raised by artificial neural development. In this workshop, we will explore existing and future approaches that aim to incorporate development into ANNs. Invited speakers will present their work with neural networks, both artificial and biological, in the context of development. Accepted submissions on contemporary work in this field will be presented and we will hold an open discussion on the topic.

Submission

A 2-4 page PDF file in LNCS format should be emailed to the workshop organizers by May 21, 2018 to be considered for acceptance in the workshop.

Accepted submissions will be distributed online and presented during the workshop.

Organizers

Julian F Miller, julian.miller@york.ac.uk
Sylvain Cussat-Blanc, sylvain.cussat-blanc@irit.fr
Dennis G Wilson, dennis.wilson@irit.fr

To contact the organizers concerning the workshop, please email devonn@irit.fr.
Details

Hosted by PPSN 2018
Workshop date: 8 or 9 September 2018
Duration/Time: half-day (to be announced)
Location: Coimbra, Portugal

About this newsletter

SIGEVOlution is the newsletter of SIGEVO, the ACM Special Interest Group on Genetic and Evolutionary Computation. To join SIGEVO, please follow this link: [WWW]

Contributing to SIGEVOlution

We solicit contributions in the following categories:

Art: Are you working with Evolutionary Art? We are always looking for nice evolutionary art for the cover page of the newsletter.

Short surveys and position papers: We invite short surveys and position papers in EC and EC related areas. We are also interested in applications of EC technologies that have solved interesting and important problems.

Software: Are you are a developer of an EC software and you wish to tell us about it? Then, send us a short summary or a short tutorial of your software.

Lost Gems: Did you read an interesting EC paper that, in your opinion, did not receive enough attention or should be rediscovered? Then send us a page about it.

Dissertations: We invite short summaries, around a page, of theses in EC-related areas that have been recently discussed and are available online.

Meetings Reports: Did you participate to an interesting EC-related event? Would you be willing to tell us about it? Then, send us a short summary, around half a page, about the event.

Forthcoming Events: If you have an EC event you wish to announce, this is the place.

News and Announcements: Is there anything you wish to announce, such as an employment vacancy? This is the place.

Letters: If you want to ask or to say something to SIGEVO members, please write us a letter!

Suggestions: If you have a suggestion about how to improve the newsletter, please send us an email.

Contributions will be reviewed by members of the newsletter board.

We accept contributions in LATEX, MS Word, and plain text.

Enquiries about submissions and contributions can be emailed to editor@sigevolution.org

All the issues of SIGEVOlution are also available online at: www.sigevolution.org

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in the Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library
- to allow users to copy and distribute the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will make every effort to refer requests for commercial use directly to you.

Editors: **Emma Hart** & **Gabriela Ochoa**
Associate Editors: **Darrell Whitley, Una-May O-Reilly, James McDermott**
Design & Layout: **Kate Simpson**