DAIS ITA - BPP18

Distributed Analytics and Information Sciences

International Technology Alliance

Biennial Program Plan 2018

15th Jan 2018 – 14th Jan 2020

Summary of Accomplishments

Written on 24th March 2020
Agreement W911NF-16-3-0001

Table of Contents

Summary of Accomplishments

Executive Summary

In this report covering the two-year BPP18 period we have chosen to highlight certain key publications for each of the 6 projects. Please refer to the Science Library http://sl.dais-ita.org/science-library for all publications during this period and the whole of the DAIS ITA program so far. At the time of writing there are a number of "in flight" publications that have not yet been published to science library since they are under review, or as yet unpublished in their external conference or journal format, however for the BPP18 period we have confirmed 36 journal papers and 182 external conference papers, all of which have been published to the science library, along with 1400+ citations for these papers so far. Our estimate is that an additional 12 journal papers and 17 external conference papers will eventually be published for the BPP18 period. The results of each project are summarised below, largely through choosing highlights from the year. In many cases these correlate with what was presented to the peer reviewers in the 2019 Annual Fall Meeting, resulting in a very good result from the peer reviewers, commending the programme on progress, publications and high-risk, high-reward scientific research. In addition to publications, a number of projects have also published various code and data assets to open source repositories as reported below. This is a good set of materials that have been delivered from a scientific and practical perspective and brings the highly productive BPP18 research period to a strong finish.

Project 1: Software Defined Coalitions

In year one a team of researchers at Yale University, IBM U.S., U.S. Naval Research Laboratory, and Imperial College investigated how to optimally set the synchronization levels between distributed SDC (Software Defined Coalition) controllers to maximize the performance of distributed set of applications in coalition. Specifically, they developed a novel algorithm using optimization and learning-based methods, which has a provable performance guarantee confirmed by extensive experimental analysis on various network topology and application scenarios. This result, at the core of SDC control plane architecture design, provides significance performance improvement for the applications running in the in future coalition networks. K. Poularakis, Q. Qin, L. Ma, S. Kompella, K. Leung, and L. Tassiulas, “Learning the Optimal Synchronization Rates in Distributed SDN Control Architectures”, to appear in IEEE Infocom 2019. https://dais-ita.org/node/2532

Researchers at Imperial College, IBM U.S., Yale University, and Defence Science and Technology Laboratory investigated the service placement issues in distributed SDC environments in coalition. A key challenge is how to improve the SDC controller scalability, especially minimizing the control-plane overhead in the placement decision making in distributed coalition networks while ensuring good performance for the applications and end users. The team developed a novel mechanism based on reinforcement learning algorithm with guaranteed performance and convergence rate. Extensive experimental analysis shows that the algorithm consistently outperforms the state-of-the-art methods in various network topologies, and also reveals important insights into how the distributed services shall be planned and placed in future coalition environments. Liang Ma, Spike Zhang, Kin Leung, Leandros Tassiulas, and Jeremy Tucker, “Q-placement: Reinforcement-learning-based Service Placement in SDC”, in IEEE ICDCS 2018. https://dais-ita.org/node/2376

A team of researchers from Yale University and IBM U.S. investigated a key issue in SDC control in dynamic coalition networks, specifically how to support frequent and/or bursty

changes in SDC resources. They created a new theoretic foundation and systematic framework for SDC control-plane mechanism that can enable continuous and fast updates of resource availability across coalition boundaries in dynamic environments. The method is based on a novel concept of update algebra, which models the control-plan operations by updates in a set-theoretic projection space, hence enabling composition of multiple changes into smaller number of updates while guaranteeing the correctness. Extensive experimental analysis shows that this new method can achieve a significant gain in the performance (16x reduction in update latency) with more efficient use of the control-plane resources (reduction in control-plane bandwidth by 30%). Geng Li, Franck Le, Yeonsup Lim and Y. Richard Yang, "Update Algebra: Toward Continuous, Non-Blocking Composition of Network Updates in SDN", to appear in IEEE Infocom 2019. https://dais-ita.org/node/3284

During year 2 researchers at Imperial, IBM U.S. and Yale developed new Deep Reinforcement Learning (DRL) methods for automated design of controller synchronization policy in Software Defined Coalition (SDC) networks. In particular, we investigated a scenario where the domain-wise communication path between two end nodes in the SDC forms a line network. We developed DRL-based algorithms to learn the synchronization policies between the source and other domains where data packets traverse. Z. Zhang, L. Ma, K. Poularakis, K.K. Leung and L. Wu, "DQ Scheduler: Deep Reinforcement Learning Based Controller Synchronization in Distributed SDN," presented and received the Best Paper Award at IEEE ICC 2019. https://dais-ita.org/node/3115 Team members at Yale and Imperial devised deep-learning models for data-plane switches at the network edge. We proposed Binarized Neural Networks (BNNs) to run on the resource-limited switches, which can significantly expedite the inference process. We further adopted a Federated Learning framework such that knowledge from different switches can be combined to shorten training and enhance intelligent decisions. Experiments show the scalability and performance benefits of the proposed design. Q. Qin, K. Poularakis, L. Tassiulas and K.K. Leung, "Line-Speed and Scalable Intrusion Detection at the Network Edge via Federated Learning," submitted to IFIP Networking 2020. https://dais-ita.org/node/4678 Researchers from Yale and IBM U.S. developed the first, systematic formulation of the software-defined internetworking (SDI) model, in which SDC exposes a programmable interface to allow coalition partners to define the inter-domain routes, while maximizing network autonomy to maintain partners' control of inter-domain export policies. They developed a blackbox optimization algorithm that finds a near-optimal policy-compliant end-to-end route. Extensive experiments show that in an SDC with over 60,000 units, the algorithm finds the optimal route within 33 route samples. Q. Xiang, J. Zhang, Y.-S. Lim, F. Le, G. Li and Y.R. Yang, "Toward Optimal Software-Defined Interdomain Routing," to appear in IEEE INFOCOM 2020. https://dais-ita.org/node/4680 Team members at Yale investigated how to integrate SDC functions into an SDN programming framework, and proposed Trident, a novel unified SDC programming framework. Trident extends algorithmic SDN programming with new primitives (e.g., stream attribute, live variable and route algebra) to reduce complexities of integrating changing states of SDC functions into SDN. Evaluations show that Trident causes very little overhead for integrating existing network management tools and SDC functions, and can handle up to O(10^5) routing requests per second with O(100) milliseconds latency. K. Gao, T. Nojima and Y.R. Yang, "Trident: Toward Distributed Reactive SDN Programming with Consistent Updates," to appear in IEEE JSAC. https://dais-ita.org/node/4681

Project 2: Generative Policy Models for Coalitions

In year one researchers at Imperial College and Purdue University came up with a novel notion and formal definition of Answer Set Grammars (ASGs), which allow annotations on

production rules written in the language of Answer Set Programming (ASP) that can express context-sensitive constraints. They showed that the complexity of the decision problems can be lowered by the use of ASGs while allowing for context-sensitive languages to be included. They also proposed a learning-based method for generating the annotations of an ASG, providing a formal framework and foundation for generative policies. M. Law, A. Russon, E. Bertino, K. Broda, and J. Lobo, "Representing and Learning Grammars in Answer Set Programming", in AAAI 2019. https://dais-ita.org/node/2512 A team of researchers from IBM U.K, IBM U.S., Imperial College, and Purdue University investigated how well generative policy models perform in real-world applications. Specifically, they evaluated generative policy models, based on machine learning (ML) approaches and ASG approach, in a connected and autonomous vehicle operation scenario, through data-driven analysis of both simple and complex tasks. They found ASG based techniques outperform shallow ML approaches for learning complex knowledge when a large number of examples are available for coping with the increase in the size of the hypothesis space. They also publicly released the annotated dataset for the considered scenario (https://github.com/dais-ita/cav_annotated_policies). D. Cunnington, I. Manotas, M. Law, G. de Mel, S. Calo, E. Bertino, and A. Russon, "An Answer Set Grammar-based Generative Policy Model for Connected and Autonomous Vehicles", submitted to IEEE ICDCS 2019. https://dais-ita.org/node/3264

During year two project 2 has designed a new symbolic learner, FastLAS, for learning policies form positive and negative examples. However often many alternative hypotheses can explain the examples, and most systems employ a bias towards shorter solutions, based on Occam's razor. Choosing the shortest hypothesis often leads to very general hypotheses being learned from relatively few examples. While this can be an advantage of symbolic learning over other machine learning approaches that need larger quantities of data, learning such general rules without large quantities of data to justify them may not be desirable in every application domain. For example, in access control, wrongly allowing access to a resource may be far more dangerous than wrongly denying access. So, learning a more general hypothesis, representing a more permissive policy, would be more dangerous than a specific hypothesis, representing a more conservative policy. Equally, for access control where the need for resources is time critical, wrongly denying access could be more dangerous than wrongly allowing access. When learning such policies, it would be useful to specify whether the search should be biased towards more or less general hypotheses. FastLAS addresses this problem by the introduction of (domain specific) scoring functions. The idea is that if there are alternative hypotheses that explain the examples, the hypothesis with the lowest score is preferred. FastLAS has also introduced a new approach for computing a smaller subset of the hypothesis space guaranteed to contain at least one optimal solution with respect to the given scoring function. This smaller search space can be orders of magnitude smaller than the full hypothesis space thus making FastLAS highly scalable.

We have developed a novel approach for learning attribute-based policies from logs and other data sources, such as directories. The approach combines in a novel way several techniques: association rule mining, clustering, statistical analysis, classification, and deep learning. The experimental evaluations, carried out on real and synthetic datasets, show that the approach is more efficient than other approaches.

We have investigated policy-based federated learning and introduced the concept of policy-based learning ensembles, in which ensemble learning is combined with policies that are generated to select a limited subset of models to be used in different situations. This dynamic

selection of models can be used to obtain a better approximation to ground truth in situations where the data used to train individual models is widely different.

Project 3: Agile Composition for Coalition Environments

In year one researchers at University of Massachusetts, Amherst; Raytheon BBN Technologies, Army Research Laboratory, and Imperial College investigated a fundamental issue of characterizing the performance trade-offs in providing analytics services to users in tactical military environments. Specifically, they developed an analytic model for allocating requesters of analytics tasks to resources on servers in tactical wireless environments, and derived the theoretical capacity, based on queueing models of the services and requests, of the distributed analytics services in serving the requests from distributed use, under various request allocation strategies. These results, validated numerically in 1-dimensional topology and being extended to more general cases, provide a fundamental insights into what the overall capacity of the distributed analytics services would be for the users, and can lay the theoretical foundation for planning and scheduling the distributed services in tactical environments. Nitish K. Panigrahyy, Prithwish Basu, Don Towsley, Ananthram Swami, Kevin S. Chan, Kin K. Leung, "A queueing theoretic model for resource allocation in one dimensional distributed analytics network", ACM SIGMETRICS Performance Evaluation Review. 46. 27-29. https://dais-ita.org/node/2384 A team of researchers from Penn State University, IBM U.S., and Army Research Laboratory, and University College London investigated the problem of measuring the states of distributed services in coalition environment, and developed methods that can infer the structure and the state of the logical network formed by network service instances through external observations of the end-to-end performance measurements. Besides the novel application of network tomography in the service network inference problem, this approach fundamentally advances the state of the art on topology discovery and network state estimation by considering general topologies with general measurement paths. Yilei Lin, Ting He, Shiqiang Wang, Kevin Chan, and Stephen Pasteris, "Looking Glass of NFV: Inferring the Structure and State of NFV Network from External Observations", to appear in IEEE Infocom 2019. https://dais-ita.org/node/2508 A team of researchers from Penn State University, IBM U.S., and Army Research Laboratory explored the problem of summarizing large datasets to be used in machine learning models in distributed analytics environments. Specifically, the team investigated the problem of constructing a universal coreset, which is a compact summary of the original dataset in the form of a small weighted set in the same data sample space, while providing provable guarantee of the machine learning models' performance. Through a thorough theoretical analysis, they overcame the limitation of the existing algorithms tailor-made to specific machine learning problems and showed that a coreset construction algorithm based on k-clustering gives provably good approximation for a broad range of ML problems. The analytical results are verified through a diverse set of data and ML models, laying the foundation for universally applicable data summarization methods in distributed coalition environments where exchanging the large dataset is prohibitive. Hanlin Lu, Ming-Ju Li, Ting He, Shiqiang Wang, Vijay Narayanan, and Kevin Chan, "A Universally Good Coreset for Distributed Machine Learning", submitted to Mobihoc 2019. https://dais-ita.org/node/3307

During year two team members at Imperial, UMass, BBN, and ARL devised a framework using Cooperative Game Theory (CGT) to share resources among mobile edge clouds (e.g., owned by coalition partners). For a monotonic, non-decreasing utility function, the game was proven to be canonical and convex. We proposed two efficient algorithms that provide Pareto optimal resource allocations and stability of the grand coalition. Experimental results validated our proposed framework. F. Zafari, K.K. Leung, D. Towsley, P. Basu, A. Swami

and J. Li, “Let’s share: A Game-Theoretic Framework for Resource Sharing in Mobile Edge Clouds,” submitted to IEEE TNSM. https://dais-ita.org/node/4664 Researchers from UMass, BBN, ARL, and Imperial extended our optimal resource-assignment method for one-dimensional to two-dimensional distributed service networks. Specifically, we developed a heuristic algorithm, which leverages the optimal dynamic-programming scheme for one-dimensional inputs to obtain approximate solutions for the two-dimensional scenario, to empirically yield request distances within a constant factor of the optimal solution. We obtain closed-form expressions for evaluating the expected communication cost. We also compared performance of various allocation policies in terms of the variance of request distance. N.K. Panigrahy, P. Basu, P. Nain, D. Towsley, A. Swami, K.S. Chan and K.K. Leung, “Resource Allocation in One-dimensional Distributed Service Networks and its Applications,” submitted to Performance Evaluation. https://dais-ita.org/node/4695 Researchers from PSU, IBM US, ARL, and UCL extended year 1’s work on robust coresets to include joint cardinality reduction (CR), dimensionality reduction (DR), and quantization, to support distributed machine learning at reduced complexity and communication cost. On combining CR with quantization, the team theoretically analysed the ML error bounds, based on which an optimization problem to minimize the ML error under a fixed budget of communication cost is formulated and efficient algorithms are developed. On combining CR with DR, the team proved several theoretical characteristics of the complexity and communication cost. All findings were validated using experiments based on real datasets. H. Lu, T. He, S. Wang, C. Liu, M. Mahdavi, V. Narayanan, K. Chan and S. Pasteris, “Communication-efficient k-Means for Edge-based Machine Learning,” submitted to ICDCS 2020. https://dais-ita.org/node/4677 Team members at UCL, IBM US, PSU have studied an online learning version of the facility location problem, which includes service placement problems but is broader than service placement alone. A novel and efficient algorithm for this problem has been designed. The team has proven that the loss of the proposed algorithm is bounded, up to a logarithmic factor, to the bound of the Hedge algorithm with an exponential action space. Furthermore, the proposed algorithm is efficient as its space and time complexity is quasi-linear. S. Pasteris, T. He, F. Vitale, S. Wang, M. Herbster, “Online Learning of Facility Locations with Theoretical Performance Guarantees,” to be submitted to NeurIPS 2020. https://dais-ita.org/node/4553

Project 4: Instinctive Analytics in a Coalition Environment

In year one a team of researchers from IBM UK and IBM US investigated advanced techniques for workflow representation which resulted in a paper at COLLA 2018 (Conference on Advanced Collaborative Networks, Systems and Applications. This paper introduced Vector Symbolic Architectures (VSAs) that are distributed representations that combine random patterns, representing atomic symbols across a hyper-dimensional vector space, into new symbolic vector representations that semantically represent the component vectors and their relationships. In this paper, the researchers extended the VSA approach and applied it to decentralized workflows, capable of executing distributed compute nodes and their interdependencies. The paper won the ‘Best Paper Award’ with invitations to submit an extended article version to one of the IARIA Journals. Simpkin, Christopher, Taylor, Ian, Bent, Graham, de Mel, Geeth, and Ganti, Raghu.“A Scalable Vector Symbolic Architecture Approach for Decentralized Workflows.” Presented at: The Eighth International Conference on Advanced Collaborative Networks, Systems and Applications COLLA 2018, Venice, Italy, 24-28 June, 2018. https://dais-ita.org/node/2397 A team of researchers at IBM UK, Penn State, Yale, Cardiff, IBM US, ARL and Dstl investigated the automatic composition of services to dynamically match operational tasks to information requirements, accounting for impact, in a many-to-many temporally and spatially complicated and complex situations. In

dynamic and agile environments, such as coalition environments, the state of the network and resources cannot be completely known in advance nor controlled due to the evolving nature of the network and constraints that may preclude partners from accessing complete state information about different parts of the system. In addition, there may be requests made to the system that have not been made before, requiring services to be created on the fly. Motivated by these observations, in this paper, we present a critical analysis of gaps in the state-of-the-art and our vision to address those through novel theoretical contributions. We envision that such formalized and theorized fundamentals will enable service elements to automatically configure themselves to perform analytic tasks based on user specified goals by taking account of context—be it system or user context. Bent, Graham, Geeth de Mel, Raghu Ganti, Tom La Porta, Gavin Pearson, Tien Pham, Sebastian Stein, Leandros Tassiulas, and Ian Taylor. "Learning Service Semantics for Self-Organization in Distributed Environments: Concepts and Research Directions." In MILCOM 2018-2018 IEEE Military Communications Conference (MILCOM), pp. 1080-1085. IEEE, 2018. https://dais-ita.org/node/2400

During year two team members in project 4, including IBM US, IBM UK, Purdue, PSU, Yale, Cardiff, and Southampton collaboratively worked together to address the problem of performing distributed analytics in a coalition environment by automatically composing complex services to dynamically match operations tasks to information resources, accounting or impact, in many-to-many temporally and spatially complicated and complex situation. A key component was to develop resource allocation for dynamically formed distributed analytics services, which was proven to be NP-hard and algorithms with near-optimal solutions were developed. This work was published in top conferences such as ICDCS 2018 and Infocom 2019. Extensions to resource allocation while taking request routing into account and truthfulness as well have been developed. The problems were formulated and shown to be NP-hard and near-optimal approaches have been proposed. These approaches were published in top conferences such as Infocom 2019 and AAMAS 2020. Another aspect of being able to create instinctive analytics is the ability to compose services on the fly, as needed by the tasks at hand. A key contribution from the team has been to develop encoding mechanisms using Vector Symbolic Architectures for composition of services - this work has been published at various venues such as FGCS 2019 and COLLA 2018. It won the best paper award at COLLA 2018. Model ranking, a method to rank models to more efficiently compose services was also developed during this period and published at Milcom 2018. Further, a key aspect of being able to compose machine learning, specifically deep learning services efficiently is energy. Novel spiking architectures for distributed execution of machine learning tasks was developed as part of project 4 and published in Nature 2019.

Project 5: Anticipatory Situational Understanding for Coalitions

In year one researchers on project 5 (ARL and Ozyegin University) in collaboration with the wider research community (Bosch Center for AI) have investigated novel techniques to provide uncertainty awareness to deep neural networks. Specifically, regarding the standard approach to training deterministic neural nets - the network is trained to minimize a prediction loss so resultant model remains ignorant to its prediction confidence. Orthogonally to Bayesian neural nets that indirectly infer prediction uncertainty through weight uncertainties, we propose explicit modeling of the same using the theory of subjective logic. By placing a Dirichlet prior on the softmax output, they treat predictions of a neural net as subjective opinions and learn the function that collects the evidence leading to these opinions by a deterministic neural net from data. The resultant predictor for a multi-class classification problem is another Dirichlet distribution whose parameters are set by the continuous output of a neural net. Sensoy, Murat, Lance Kaplan, and Melih Kandemir. "Evidential deep

learning to quantify classification uncertainty." In Advances in Neural Information Processing Systems, pp. 3183-3193. 2018. https://dais-ita.org/node/2574 Researchers at IBM UK, Cardiff University and IBM UK further developed the interpretability aspects of their research, focusing on role-based interpretation and how the role of the user is important in determining appropriate explanations. More specifically they assert that a machine learning system's interpretability should be defined in relation to a specific agent or task. We should not ask if the system is interpretable, but to whom is it interpretable. This research identifies different roles that agents can fulfil in relation to a machine learning system to help researchers, system developers, and regulatory bodies structure their reasoning about interpretability requirements in machine learning systems. Tomsett, R.; Braines, D.; Harborne, D.; Preece, A.; and Chakraborty, S. 2018. Interpretable to whom? A role-based model for analysing interpretable machine learning systems. In 2018 ICML Workshop on Human Interpretability in Machine Learning (WHI 2018). https://dais-ita.org/node/2722

For project 5 task 1, the second year of BPP 2018 has seen significant success in two strands of research. In the first, Cardiff and ARL advanced the state-of-the-art in uncertainty-aware learning and reasoning with papers at AAAI2019 and AAAI2020, and two journal papers in preparation. This work addresses the coalition problem of providing robust learning and inference in the presence of sparse data and complex dynamic situations. Via a probabilistic logic programming approach that can reason in the presence of uncertain probabilities, we combine state-of-the-art performance with the flexibility afforded by logic programming. Thus, we provided efficient mechanisms for enabling logical, symbolic reasoning both with epistemic and aleatoric uncertainty. In the second strand, led by UCLA and Cardiff, we advanced the state-of-the-art in complex event processing (1) fusing proxy models; (2) using hybrid neuro-symbolic approaches; and (3) using purely neural approaches yet enabling humans to provide regularizing rules. We overcame the need for specialized learnt models by using human expertise in guiding the fusion of off-the-shelf systems. Our novel work in neural architectures for complex event processing in an edge, coalition context led to a well-received demonstration at AFM2019 and a paper subsequently submitted to IJCAI2020. Achievements in project five task 2 addressed the coalition need for ML services that are more explainable and more robust to adversarial attacks (for assurance); thus more exploitable in human-machine teams. Work by UCLA and IBM US on using gradient-free optimization for generating adversarial examples in black box setting - published at GECCO2019 - effectively opens up the space of adversarial attacks by designing a query-efficient black box attack based on gradient-free search strategy. This expands the scope of adversarial attacks even to networks with non-differentiable operators. IBM US led work on analysis of model poisoning as an attack vector in a federated learning setting, published at ICML2019. We demonstrated that the high-dimensional parameter space of neural network model can be effectively perturbed to introduce targeted backdoor into the global model. Research on interpretability led by IBM UK and Cardiff, published at AAAI2020, investigated the inconsistency of perturbation-based metrics in ranking interpretability mechanisms. We showed that it is impossible to perturb the input pixels and not generate artefacts that are in turn used by the model for its classification result: perturbation of the input pixels always introduces bias in the model score and hence metrics that rely on this score introduce bias in their ranking.

Project 6: Evolution of Complex Adaptive Human Systems

In year one researchers from Cardiff University and Yale have published substantial work as a Nature Scientific Reports journal paper which has gained significant coverage in the UK mainstream media. In this work the researchers observe that prejudicial attitudes are widely

seen between human groups, with significant consequences. Actions taken in light of prejudice result in discrimination and can contribute to societal division and hostile behaviours. They define a new class of group, the prejudicial group, with membership based on a common prejudicial attitude towards the outgroup. They assume that prejudice acts as a phenotypic tag, enabling groups to form and identify themselves on this basis. Using computational simulation, they then study the evolution of prejudicial groups, adopting indirect reciprocity as the social dilemma, observing how cooperation and prejudice coevolve, with cooperation being directed in-group. Whitaker, Roger M., Colombo, Gualtiero B., and Rand, David G. "Indirect reciprocity and the evolution of prejudicial groups." Scientific reports 8, no. 1 (2018): 13247. https://dais-ita.org/node/2356 Researchers from UMass, IBM UK and Cardiff University extend traditional network motif analysis to consider directed edges and temporality. They propose a new methodology that utilizes feature representation for network classification based on the temporal motif distribution of the network and a null model for comparing against random graphs. Experimental results show an improved accuracy by up 10% compared to the state-of-the-art embedding method in network classification, for tasks such as classifying network type, identifying communities in email exchange network, and identifying users given smartphone app-switching behaviours. Tu, Kun, Li, Jian, Towsley, Don, Braines, Dave, Turner, Liam D. “Network classification in temporal networks using motifs.” In Proc. of the 3rd ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data (2018). https://dais-ita.org/node/2496

Year two task 1 highlights include:

• Cultural drift through organisations and networks. Culture represents ideas and beliefs that are transmitted through groups by social interaction. Understanding cultural spread represents valuable intelligence capability. We have developed a new approach allowing “peer pressure” to be modelled https://dais-ita.org/node/3250. This generalises the state-of-the-art in cultural modelling. Real structures such as coalitions, are now being examined in BPP20.

• Understanding the evolution of identity fusion. Personal identity, or how individuals choose to be represented by what's important to them in world, is a psychological concept linked to the "devoted actor", a leading explanation for extremism. Extending our work https://dais-ita.org/node/2356 on group formation, we have introduced a new quantitative model for identity, using traits to represent both individuals and groups. We have successfully used this to explore stereotyping https://dais-ita.org/node/3433. We have further examined the role of shared identity on cooperation https://dais-ita.org/node/4683.

• Cognitive dissonance and group evolution. Cognitive dissonance is the friction that results from inconsistency between one’s own belief and actions, and those of others. This motivates personal behaviour that can affect how relations are held with others. We have formulated a comprehensive model to examine how cognitive dissonance coevolves with network structure https://dais-ita.org/node/4682 . Separate to the ITA, interest from Dstl has led to a comprehensive review by Dstl on the psychological relevance of the concepts https://dais-ita.org/node/3433

Year two task 2 highlights include:

• Comparisons of Metrics and Models of Multiple Social Networks. Through application of motif analysis to diverse social networks, it has been possible to determine the scope for categorization of networks based on their real-world behavioural context. Extension of motif analysis to include tetrads, has also been undertaken, showing that additional value can derived from consideration of larger substructures. The capacity to make location-based prediction, using features of the network, has also been addressed https://dais-ita.org/node/3746 .

• Dynamic Motifs in Dynamic Networks. We have explored the value of motifs as predictive features for network classification for static and dynamic networks https://dais-ita.org/node/3435 . We improve upon the state-of-the-art by accounting for presence of directed edges and temporal changes. A novel network embedding methodology has been developed that goes beyond node representations for network classification tasks. The approach shows that motifs provide useful additional features for prediction.
• Motifs and Emergence of Inter-group Behaviour. Work has focused on the interplay between groups of conflicting authors who interact online. This has shown how motifs can be used as a predictive feature that allows controversial and non-controversial social media to be distinguished https://dais-ita.org/node/3388 . The approach is particularly novel as it avoids semantic analysis and is transferable to other forms of social media. Furthermore, motifs have played a fundamental role in assessing patterns of individual-level human behaviour, expressed through interactions with technology. This has been published in the Royal Society Open Science Journal https://dais-ita.org/node/3251 .
• Understanding the role of sub-structures for network intelligence. Motif analysis is predicated on assessing the under or over representation of sub-structures, such as triads within a network. The under or over representation of triads, given their potential role in the network, provides further useful context. We are now exploring application of this framework to further understand the agency of actors in the evolution of terrorist networks as recently examined using exponential random graph modelling https://dais-ita.org/node/3719 .

Additional details

Additional details about the BPP18 period plans and goals can be found at https://dais-ita.org/BPP18 A summary of accomplishments and publications for the BPP18 and the whole of the program can be found at https://dais-ita.org/project_qprs and also https://dais-ita.org/qpr_archives

In the above summary of accomplishments, we have chosen to highlight certain key publications for each of the 6 projects. The next section of this report is a summary of publications occurring during the two year BPP18 period. Please refer to the Science Library (http://sl.dais-ita.org/science-library) for all publications during this period and the whole of the DAIS ITA program so far.

Summary of Publications

Introduction

The data and statistics in this document are summarised from the publicly available DAIS ITA Science Library [http://sl.dais-ita.org/science-library]. In some cases there are publications listed in this document that are not yet published to Science Library, e.g. because they have not yet been published in their external conference or journal venue (See Appendix 1). Eventually, all successfully accepted papers will be published to Science Library, even after the BPP18 period has concluded. The papers counted in this summary document will tally closely to those reported throughout the BPP18 period in the Quarterly Progress Reports (QPRs) but they are reported separately here as a single succinct summary as the repeating periodic nature of the QPR document means that papers are regularly reported in numerous quarters as they progress from submitted to accepted and published, plus each QPR only deals with the papers relevant to that quarter. The purpose of this stand-alone summary document is therefore to provide a short and simple overall summary for the period, backed up by detailed statistics on science library and more detailed contextual reporting in each of the QPRs.

Overall summary

The BPP18 period (24 months) saw 36 journal papers published at external peer reviewed journals, and 182 conference papers published at peer reviewed external conferences or workshops. This compares to 6 journal and 90 conference/workshop papers in the IPP period (16 months). This in broken down by project in the table below, with citation data[1] for each project and period. The IPP data is included purely for comparative purposes.

Period		Journal	External Conference	Citations
BPP18 Jan-18 to Jan-20	BPP18 P1	13 (0)	28 (2)	459
	BPP18 P2	4 (0)	31 (4)	105
	BPP18 P3	4 (0)	25 (5)	267
	BPP18 P4	9 (0)	25 (6)	100
	BPP18 P5	3 (0)	47 (5)	353
	BPP18 P6	3 (0)	16 (1)	47
	Cross-project	0	10	70
	Total	**36**	**182**	**1491**
IPP Sep 2016 - Jan-2018	IPP P1	1 (0)	10 (2)	80
	IPP P2	0 (0)	23 (6)	148
	IPP P3	3 (0)	8 (3)	299
	IPP P4	0 (0)	13 (2)	31
	IPP P5	0 (0)	7 (6)	42
	IPP P6	2 (0)	18 (5)	297
	Cross-project	0	11	44
	Total	**6**	**90**	**941**

[1] Citation data was taken manually from Google Scholar for each DAIS paper on 17th Feb 2020.

Table 1: Overall confirmed external publication summary (by project)

Period		Journal	External Conference	Patent	Citations
BPP18 Jan-18 to Jan-20	BPP18 P1	14 (0) *[+1]*	31 (2) *[+3]*	0	464 *[+5]*
	BPP18 P2	5 (0) *[+1]*	35 (4) *[+4]*	0	110 *[+5]*
	BPP18 P3	4 (0) *[+0]*	29 (5) *[+4]*	1 (0) *[+1]*	271 *[+4]*
	BPP18 P4	19 (0) *[+10]*	29 (6) *[+4]*	2 (0) *[+2]*	164 *[+64]*
	BPP18 P5	3 (0) *[+0]*	49 (5) *[+2]*	0	353 *[+0]*
	BPP18 P6	3 (0) *[+0]*	16 (1) *[+0]*	0	47 *[+0]*
	Cross-project	0	10	0	70 *[+0]*
	Total	**48** *[+12]*	**199** *[+17]*	**3** *[+3]*	**1659**
IPP Sep 2016 - Jan-2018	IPP P1	1 (0) *[+0]*	11 (2) *[+1]*	0	80
	IPP P2	0 (0) *[+0]*	23 (6) *[+0]*	0	148
	IPP P3	3 (0) *[+0]*	9 (3) *[+1]*	0	299
	IPP P4	0 (0) *[+0]*	13 (2) *[+0]*	0	31
	IPP P5	0 (0) *[+0]*	7 (6) *[+0]*	0	42
	IPP P6	2 (0) *[+0]*	18 (5) *[+0]*	0	297
	Cross-project	0	11	0	44 *[+0]*
	Total	**6** *[+0]*	**92** *[+2]*	**0**	**941**

Table 2: Possible external publication summary (by project), including papers not yet confirmed as published, and found in Google Scholar but not reported in CENSE

Breakdown by organisation

The science library is able to show the contribution from each organisation but is not currently able to break this down by BPP period. Also, the additional 55 as-yet-unpublished papers shown earlier in table 2 are not included in these organisation statistics as they are generated directly from science library, and those additional papers are not yet published to science library. The figures shown in blue in table 3 below are computed values and may be useful for comparative purposes.

Organisation	External papers	# authors	Citations	papers / author	citations / author
Airbus	4 *[+0]*	6 *[+0]*	16 *[+0]*	0.67	2.67
BAE	10 *[+0]*	5 *[+0]*	96 *[+0]*	2.00	19.20
IBM UK	99 *[+2]*	19 *[+0]*	297 *[+0]*	5.21	15.63
IBM US	180 *[+15]*	53 *[+7]*	1168 *[+20]*	3.40	22.04
Raytheon BBN	7 *[+1]*	3 *[+0]*	3 *[+2]*	2.33	0.43
Cardiff	85 *[+3]*	26 *[+1]*	324 *[+0]*	3.27	12.46
Imperial	57 *[+3]*	14 *[+0]*	609 *[+3]*	4.07	43.50
PSU	41 *[+1]*	16 *[+2]*	305 *[+3]*	2.56	19.06

Purdue	62 *[+15]*	24 *[+9]*	255 *[+65]*	2.58	10.63
Southampton	22 *[+1]*	12 *[+0]*	62 *[+2]*	1.83	5.17
Stanford	2 *[+0]*	3 *[+0]*	6 *[+0]*	0.67	2.00
UCL	24 *[+2]*	13 *[+0]*	157 *[+0]*	1.85	12.08
UCLA	25 *[+1]*	23 *[+0]*	481 *[+1]*	1.09	20.91
UMass	24 *[+2]*	8 *[+1]*	119 *[+2]*	3.00	14.88
Yale	54 *[+2]*	22 *[+1]*	189 *[+0]*	2.45	8.59
ARL	82 *[+0]*	19 *[+0]*	510 *[+0]*	4.10	25.75
Dstl	40 *[+1]*	11 *[+0]*	142 *[+0]*	3.55	12.91

Table 3: External paper (journal and conference) publication statistics for each DAIS organisation (for IPP and BPP18 combined)

Yearly statistics

Another view on the science library is the statistics view[2]. The data from this view is recreated below and split across two tables for each of layout. As was the case for the previous section, the additional 55 as-yet-unpublished papers shown earlier in table 2 are not included in these statistics as they are generated directly from science library, and those additional papers are not yet published to science library. Also, whilst these statistics are not currently capable of breaking down into different IPP/BPP periods, the BPP18 period (15th Jan 2018 – 14th Jan 2020) relates approximately to the years 2018 and 2019 in these tables.

Year	Total External Publications	Single Institute Paper Count		Collaborative Paper Count		International Paper Count		Government Paper Count	
2016	4 *[+0]*	1 *[+0]*	25%	3 *[+0]*	75%	1 *[+0]*	25%	1 *[+0]*	25%
2017	95 *[+2]*	10 *[+0]*	11%	85 *[+2]*	89%	62 *[+0]*	65%	28 *[+0]*	30%
2018	110 *[+1]*	9 *[+1]*	8%	101 *[+0]*	92%	54 *[+0]*	49%	37 *[+0]*	34%
2019	131 *[+27]*	29 *[+15]*	22%	102 *[+12]*	78%	70 *[+7]*	53%	40 *[+3]*	31%
2020	8 *[+4]*	4 *[+4]*	50%	4 *[+0]*	50%	3 *[+0]*	38%	2 *[+0]*	25%
Total	348 *[+34]*	53 *[+20]*	15%	295 *[+14]*	85%	190 *[+7]*	55%	108 *[+3]*	31%

Table 4: Collaborative publishing behaviour

On DAIS ITA collaborative publications are encouraged, especially international ones, and in particular those including government authors. We see a low percentage of single institute publications (15%), and of the 85% collaborative publications, 55% of the total papers are international collaborations, and 31% of all publications involve government co-authors. This are improved over the IPP period, but started strongly even in IPP.

[2] Available at http://sl.dais-ita.org/science-library/statistics

Year	Journal Paper Count		External Conference Paper Count		Patent Count		Active ITA Authors	Total Citations
2016	0 *[+0]*	0%	4 *[+0]*	100%	0 *[+0]*	0%	9 *[+0]*	26 *[+0]*
2017	7 *[+0]*	8%	88 *[+2]*	92%	0 *[+0]*	0%	134 *[+3]*	867 *[+15]*
2018	17 *[+0]*	16%	93 *[+1]*	84%	0 *[+0]*	0%	148 *[+0]*	696 *[+3]*
2019	28 *[+10]*	21%	102 *[+16]*	78%	1 *[+1]*	1%	190 *[+23]*	528 *[+55]*
2020	2 *[+2]*	25%	4 *[+0]*	50%	2 *[+2]*	25%	22 *[+11]*	20 *[+20]*
Total	54 *[+12]*	13%	291 *[+19]*	87%	3 *[+3]*	1%	278 *[+28]*	2317 *[+93]*

Table 5: Publication by external type, authorship and citations

The data in table 5 is largely covered already in the earlier parts of this document, but it is useful to see how journal publications ramps up from 2017 to 2019. It is also good to see the rise in active authors, including those outside the 16 core DAIS organisations. The citation data shows how many citations have been gained for papers written in each year. It is not surprising to see that the earlier years have higher citation totals, mainly because a longer period of time has passed in order to accrue citations.

Finally, not shown in these tables, but available directly from the underlying knowledge graph: The total number of organisations that have co-authored DAIS papers on science library is 78 organisations, from 18 countries.

Annual Fall Meeting publications

In addition to the external publications covered in this document there are also numerous internal publications, technical reports and other documents already published to science library:

- 81 internal conference papers (AFM papers)
- 5 technical reports
- 52 other documents (demos, posters, PhD theses, book chapters, invited talks etc).

The analysis of the unpublished external papers has also identified 10 AFM 2017 papers, 46 AFM 2018 papers, and 34 AFM 2019 papers that can also be published to science library, assuming the external publication to which they relate has already been published. This will be the case for all of the AFM 2017 papers, many of the AFM 2018 and some of the AFM 2019 papers. As with the other publications we will assess these throughout the BPP20 period and publish periodically. Whilst these "internal" AFM papers are usually duplicates of papers published elsewhere externally, they are useful to also publish to science library since each has a corresponding poster from the AFM event which is published alongside the paper on science library.