

A Survey of Methodology in Self-Adaptive Systems Research

Barry Porter, Roberto Rodrigues Filho, Paul Dean

School of Computing and Communications

Lancaster University

Lancaster, UK

Email: {b.f.porter,r.rodriguesfilho,p.dean1}@lancaster.ac.uk

Abstract—Major research venues on autonomic and self-adaptive systems have been active for 16 years, exploring and building on the seminal vision of autonomic computing in 2003. We study the current trajectory and progress of the research field towards this vision, surveying on the *research questions* that are asked by researchers and the *methodological practice* that they employ in order to answer these questions. We survey contributions under this lens across the three main venues for primary research in autonomic and self-adaptive systems work: ICAC, SASO, and SEAMS. We examine the last three years of contributions from each venue, totaling 210 publications, to gain an understanding of the dominant current research questions and methodological practice – and what this shows us about the progress of the field. Our major findings include the general research questions focusing one level below the highest autonomy vision; methodological practice being split almost evenly between real-world experiments and simulation; a general a high positive results bias in publications; and a potential concern over low levels of repeatability across most contributions.

I. INTRODUCTION

Major research venues on autonomic and self-adaptive systems have been active for 16 years, following the seminal vision on autonomic computing presented by Kephert and Chess [1] in 2003. While the vision and end-game for autonomic computing may be clear, the current trajectory of the research field towards this vision is less obvious.

We study this topic by examining the prevailing trends in *research questions*, *methodology*, and *results* for self-adaptive systems research across the three dominant communities representing the field: ICAC, SASO, and SEAMS. This study offers cross-community insight into the questions that researchers are trying to answer, and the mechanisms with which the robustness of the answers are judged.

While there is a wide range of existing survey work in self-adaptive systems, very little of this is dedicated to examining methodological practice. The majority of existing survey material is instead focused on examining the range of technical approaches used in the field – such as the seminal survey by Salehie and Tahvildari (2009) [2], which presents a taxonomy of self-adaptive systems research in terms of how adaptation itself is governed. A broad range of similar approach-focused surveys have been reported both for the self-adaptive field as a whole and various sub-fields (e.g. including [3], [4]).

Research that specifically examines *methodology* in self-adaptive systems research is far less common, with the survey

by Weyns et al. (2012) [5] being the closest to our own objectives. This survey examines the claims made in self-adaptive research papers and the level of evidence that is provided to support those claims, and uses publications from SEAMS as the inclusion criteria. Some of the analysis criteria used in this survey overlap with those used for our own methodological focus, though our study cross-cuts the three major venues rather than sampling from a single one. Our work therefore updates our understanding of current research questions and methodological practice, and also takes a broader view across all three major sub-communities for self-adaptive systems research. Our study also offers a fresh lens through which to view the progress of the research field towards the ultimate vision of autonomic computing. We discuss a comparison to the relevant subset of results from this survey in Sec. III-D.

Our specific research questions for this study are:

- 1) What kinds of question are self-adaptive systems researchers trying to answer?
- 2) What are the dominant methodologies used to answer these questions?
- 3) How close is progress towards the most advanced aims of the original autonomic computing visions [1], [6], and how may methodological practice affect this?

To answer these questions we have studied all papers in detail published at the last three years of the ICAC, SASO, and SEAMS events (210 papers in total). This approach provides an insight into current methodology and also offers a partial view of potential trends over time.

Our major findings are that (i) the vast majority of research aims at level 4 autonomy across a broad range of domains, with very little research touching on the highest level 5 autonomous systems; (ii) the majority of research targets traditional compute performance optimisation, as well as studying research tools or methods of modelling autonomy; (iii) parametric adaptation remains cumulatively the most dominant study mechanism; (iv) there is a very strong bias in the field toward positive results; and (v) repeatability of experimental results is generally poor.

In the remainder of this paper we first present our survey methodology in detail in Sec. II, then present our results in Sec. III structured towards our three research questions above. We present a closing discussion and conclusion in Sec. IV.

II. METHODOLOGY

In this section we present our inclusion / exclusion criteria for contributions that fall within the scope of our survey, and the coding methodology we use within the included papers.

In inclusion criteria, our research questions relate to general system-adaptive systems research; as such, we include all papers published at leading primary research venues. We use ICAC, SASO, and SEAMS as our focus for these venues – all three being long-established as the leading venues for the field. We do not include papers from journals, such as Transactions on Autonomous and Adaptive Systems, as a simple way to avoid influencing our results from repeated research within extended versions of conference papers.

From these three venues, we then selected the most recent three years as a representative sample of currently dominant research questions and methodological practices, which also provides some initial insight into potential trends over time. This yields **210** papers in total for inclusion in our study.

Prior to reading these papers, we developed a coding strategy with which to answer our research questions. This coding included a set of column headings and possible values for each paper in a row. We then sampled a subset of papers and extended our coding with major missing columns relating to research methodology. Our final coding columns are presented in Table I-II along with brief explanations of each.

TABLE I
SUMMARY OF CODING CATEGORIES FOR METHODOLOGY

Column / coding heading	Explanation
Methodology	What methodology is used to answer the research question? (for example simulation, formal proof)
<i>Real-world experiment</i>	The real system was built and placed into a real (or approximately real) environment for study.
<i>Simulated</i>	A simulator was used to approximate both the system and its environment.
<i>Simulated/real</i>	The work is partly built and evaluated in reality and partly in simulation.
<i>Theoretical analysis</i>	The contribution is evaluated using a theoretical framework.
<i>Systematic review</i>	The contribution is a survey.
Findings	Are the results reported by the research positive, negative, or neutral?
<i>positive/neutral/negative</i>	
Repeatable	Is the study repeatable? (source code, clear and precise methodology, etc.)
<i>yes/no</i>	
Dataset	If applicable, does the research provide a public dataset for repeatability, or is the dataset private or synthetic.
<i>public/private/synthetic/mixed</i>	
Real system built	Has a real system been built with which to study the autonomy?
Real system compared	If a real system has been built, is a real system used for comparative study?
<i>yes/no</i>	

We then read each paper and completed our coding for each one, using consistent values where possible. In this stage each researcher involved in the study was assigned an independent subset of papers. The lead author completed their assignment

TABLE II
SUMMARY OF CODING CATEGORIES FOR RESEARCH QUESTIONS

Column / coding heading	Explanation
Application domain	What is the application domain of the study? (<i>coding is self-describing</i>)
Topic	What system aspect is autonomy being used to improve?
<i>Performance</i>	Maximise compute (CPU) performance.
<i>Research tool</i>	A platform, survey, exemplar, or other tool for the community.
<i>Modeling</i>	A study or new method on how to model autonomy, including verification.
<i>MAPE-K optimisation</i>	Approach aiming to enhance elements of the MAPE-K concept, such as adaptive monitoring.
<i>Distributed coordination</i>	Study of how to best achieve coordination between multiple adaptive systems, including multi-agent systems.
<i>Other optimisation</i>	Optimise a system element not covered by another category (e.g., resource usage).
<i>Energy efficiency</i>	Minimise the energy usage of a system.
<i>Learning quality</i>	Enhance the quality of machine learning decisions that control adaptation.
<i>Fault tolerance</i>	Minimise impact of failures.
<i>Security</i>	Minimise impact of cyber security attacks.
<i>Alt</i>	Research does not fit into any above code.
Explicit RQ	Are research questions explicitly stated, or do they need to be inferred by the reader?
<i>yes/no</i>	
Autonomy Level	What level of autonomy is being attempted, using the taxonomy from [6]?
<i>Levels 1-5, see [6]</i>	Level 1: lowest; Level 5: highest autonomy.
Autonomy controller	What technical approach is used to control autonomous processes? (i.e., a hand-written protocol, a policy, machine learning)
<i>Learning</i>	Uses machine learning to decide actions.
<i>Protocol</i>	Uses a hand-crafted executable protocol with predetermined actions and contexts.
<i>Policy</i>	Uses a hand-crafted (non-executable) policy to decide actions.
<i>MAS</i>	Uses a multi-agent systems approach.
<i>Control theory</i>	Uses control theory to decide actions.
Adaptation type	What kind of adaptation is considered?
<i>Parametric</i>	Adapts parameter settings of a system.
<i>Structural</i>	Adapts the composition of modules used to form a system.
<i>Meta</i>	Adapts elements of the MAPE-K system.
Whole-system autonomy	Does the autonomy apply to an entire system, or a domain-specific sub-region
<i>yes/no</i>	

first, with the other researchers then using this coding as a template for their assignment. Following completion of the remaining assignments, the lead author undertook a normalisation phase to moderate all work and ensure consistency of coding terminology. This normalisation phase also reduced the number of coding terms used to a maximum of 12, by combining low-count coding values into higher-level coding values where needed. All of our normalised coding values are listed and briefly explained in the same tables (italicised rows).

Based on the results of coding we then draw answers to our research questions. The full raw dataset of our results, plus the normalised version of this data on which our observations are based, is made publicly available¹.

¹<https://github.com/barryfp/acsos2020survey>




Fig. 1. Our coded data on the main application domain of the contribution.

III. FINDINGS

In this section we present the findings of our research. We first focus (Sec. III-A) on research questions, exploring the topics and autonomy levels that researchers are investigating. We then examine the methodologies used (Sec. III-B) to answer these research questions, and the findings reported using these methodologies.

Because a normalisation process as described above is likely to make more unusual research less visible, we also dedicate a specific subsection (Sec. III-C) to less mainstream work in order to counter normative reinforcement in our own reporting. Lastly, we present a comparison where possible between our data and the results of the 2012 Weyns study [5] in Sec. III-D.

A. Research Questions

This section reports on the kinds of topics that researchers are most interested in. For completeness we note that we cannot distinguish here between the volume of research being submitted under these different areas, and the volume of work being accepted, as we cannot observe submitted-but-rejected papers. It is possible that these results are therefore biased by the review processes of the respective communities in selecting certain kinds of research over others. For simplicity, however, our analysis here assumes that the volume of work submitted under each area roughly correlates with the volume accepted for publication in those areas.

Fig. 1 shows the primary application domain of the set of papers included in our study. The dominant area here is *data centres*, which encompasses cloud systems and other non-containerised data center technology. As well as being the cumulative majority, this is also consistently a dominant theme in each individual year of our study. The next largest body of research over the three years in *general self-adaptive systems*, representing any general theory that transcends multiple application domains or is a theoretical contribution without a specific target domain. In third place, the *other* category aggregates a wide range of smaller application domains which each have 5 papers or less in our studied set of papers – including agriculture, social systems, healthcare, and disaster management, among others.

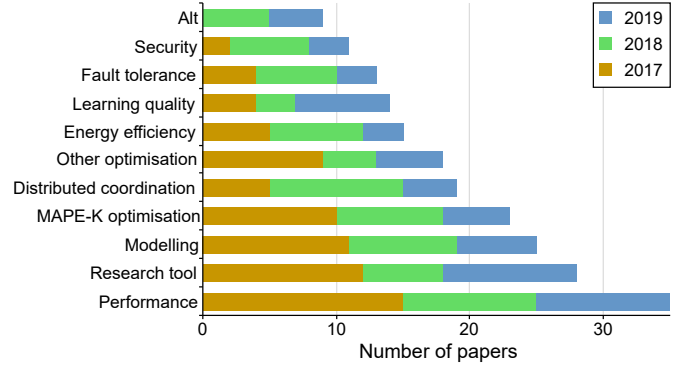


Fig. 2. Our coded data on the objective of the self-adaptive or autonomic system, such as improving performance or security.

The significant majority focus on data centres here is an interesting feature, as there seems to be no particular evidence that this domain has especially difficult or interesting autonomous control features compared to any other domain. The dominance of this application domain as a key area of study for the community may therefore be caused by alternative factors, such as relative ease of entry in equipment cost, or perceived impact probability of the research compared to other domains. On this topic, an interesting study in itself may be to gain a deeper understanding of the relative autonomy dimensions of each application domain, and therefore the level of complexity and challenge that different domains offer to researchers.

Fig. 2 shows the primary objective of autonomy or self-adaptation within the research question – for example, aiming to optimise performance, enhance security, etc. The dominant cumulative category is *performance*, where the autonomy is specifically there to enhance computational (CPU) performance such as response time. The next two most dominant categories of contribution are then *research tools* (including platforms, exemplars, and surveys) and *modelling*, which relates to different approaches for modelling autonomic behaviours (including formal verification).

The high level of contributions on research tools and modeling suggests that the research community has achieved little consensus on which tools to use to support self-adaptive systems, or how best to model that adaptation. This is anecdotally reflected in the detail of the research papers, which use a large range of specialised frameworks, languages, and platforms to support or study different aspects of self-adaptive behaviours.

If we exclude modeling approaches and research tools, the *performance* objective is followed most closely by *MAPE-K optimisation*, encompassing approaches which attempt to enhance the MAPE-K theory itself, and then *distributed coordination*, covering research which examines how multiple nodes or agents can coordinate to achieve a task. The next most dominant single-topic categories² are then self-adaptive behaviours which seek to enhance *energy efficiency*, *learning*

²Note that the category *Other optimisation* is a catch-all with no more than 5 papers from any specific category, such the use of autonomous processes to optimise resource usage or vehicle routing.




Fig. 3. Our coded data on the autonomy level involved in the research questions of each paper, using the definitions from Ganek and Corbi [6].

quality, *fault tolerance*, and finally *security* – though each of these topics is less than half the number of contributions versus that of *performance* as an objective.

This mixture of objectives suggests a healthily diverse set of research aims. If we compare this list to the original vision of autonomic computing from Kephart and Chess [1], the concepts of self-optimisation, self-healing (fault-tolerance), and self-protection (security) are all clearly present, though there is little focus on self-configuration.

Fig. 3 shows the level of autonomy that research questions are exploring, coded according to the original taxonomy by Ganek and Corbi [6], where level 1 *requires extensive highly skilled staff to analyse log data* and level 5 *is integrated components dynamically managed by business rules*. The great majority of research examines autonomy or self-adaptation at level 4, which is defined as *system monitors, correlates, and takes actions*, while human experts manage the link to higher level objectives. The next highest number of contributions is at level 3, which is defined as *system monitors, correlates, and recommends actions*. Research examining level 5 autonomous systems is very limited, with just two papers operating at level 4 but touching on level 5 by exploring how scoped high-level requirements can be translated to adaptive systems implementation and objectives [7], [8].

Overall this result raises the question of what fundamental questions remain to be understood about level 4 autonomous systems, and whether the community could begin to encourage and take on more opportunities or challenges for level 5 autonomous systems research.

Fig. 4 shows the specific kind of adaptation that research questions are exploring, coded as parametric, structural, and meta. This indicates that *parametric* is cumulatively the most popular research study over the three years of our study, with *structural* being in second place (though for 2019 it is the most dominant single-year adaptation type).

We posit that the cause of this general bias is due to relatively few real-world systems supporting hot-swapping of behaviour for structural adaptation, making researchers more likely to take an existing system and add an autonomic controller which tunes its parameters. However, parametric




Fig. 4. Our coded data on the type of adaptation being studied. *N/A* here represents papers that do not study a particular kind of adaptation, such as new platforms, surveys, or generic control theory.




Fig. 5. Our coded data on the technical approach used to implement an autonomic controller. The *N/A* category here covers papers that do not have such a controller.

adaptation alone is unable to reach the highest levels of autonomy where high-level business logic is automatically operationalised into a working system and then automatically optimised according to a desired objective. Structural adaptation is arguably able to achieve this *level 5* objective more extensively by supporting the automated composition of existing blocks of behaviour into new systems. From this perspective it seems encouraging that 2019 featured the highest volume of structural adaptation research, but we speculate that this needs to become the dominant research approach in order for the community to reach higher autonomy levels.

Fig. 5 shows the most common approaches to implementing an autonomic/adaptation controller within the studied contributions. In this coding, *policy* refers to a non-executable document which specifies how adaptation should proceed; *protocol* refers to an executable algorithm which specifies how adaptation should proceed using a set of fixed rules; and *learning* refers to some attempt to automatically learn the best strategy for adaptation. The use of learning to control adaptation is cumulatively the dominant approach, while fixed-role protocols are a close second.

We view this result as encouraging, since reaching high levels of autonomy will require systems to be able to learn for

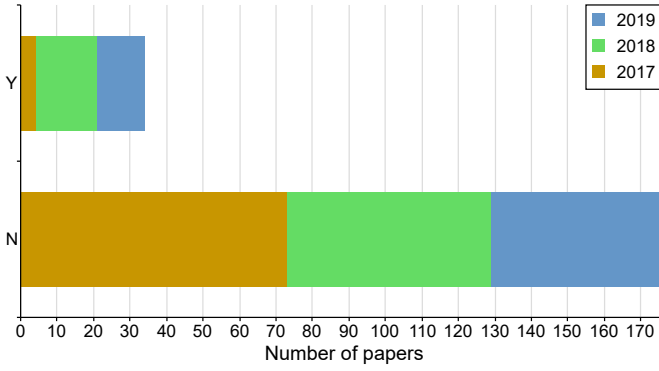


Fig. 6. Our coded data on whether or not a paper included an explicit research question to define the contribution.

themselves, rather than relying on detailed expertise and predictions of engineers encoded in policies/protocols. However, the number of contributions using fixed rule-based protocols and policies remains high, suggesting that many autonomous systems being studied do still require significant engineering expertise to generate their rules.

Fig. 6 shows the results of our coding on whether or not a paper had an explicit research question. This is slightly subjective, as it requires the coding researcher to decide whether or not the written text represents a clearly formed question, but we were cautious to give the benefit of the doubt towards ‘yes’ wherever possible. The results here are stark: the majority of research does not carry an explicit research question, which requires the reader to determine what the question is. In some cases the research question can be easily constructed by reverse-engineering elements of the motivation, for example the phrases:

“We contribute a complementary approach to tackle the problem of exhaustive analysis of large adaptation spaces using online machine learning.” [9]

“We propose a paradigm of distributed emergent software to reduce the complexity of developing these systems.” [10]

Can respectively be converted to:

“Can machine learning approaches reduce the apparent search space size for self-adaptive systems?”

“Can diverse distributed system designs be autonomously formed by real-time using emergent systems concepts?”

However, while these specific translations are relatively straightforward, in our view it is better practice to directly state the research questions of a contribution. In our experience this can help both the researchers and the readers to clearly scope the contribution of research and its context.

Fig. 7 shows the number of papers for which the research question is investigating *whole system autonomy*. We define this as any approach which encompasses every element of a system to make collective decisions about all of those elements. The counter-case is where a researcher has identified a specific part of a system to which an autonomic controller can be added and tuned, while the rest of the system does not benefit from the autonomic approach.




Fig. 7. Our results on whether a paper considers its autonomic control approach to encompass a whole system, or a sub-element of a system.

Overall, relatively few contributions take a whole-system-autonomy approach, with the great majority isolating specific parts of a system to add a targeted autonomic controller.

Examples of specialised controllers include client-side query shaping to optimise key/value store performance [11]; finding ideal CPU/GPU mixtures for Apache Spark performance tuning [12]; or making predictions of hard disk failures from standard metrics [13].

By comparison, examples of whole-system-autonomy include predictive models for cloud auto-scaling [14], product-line-based quality of service analysis which considers all features and their compositions for a system [15], or automated assembly/re-assembly of entire systems [10].

In our view, this characteristic has significant implications for reaching a higher level of autonomy: systems at level 4 autonomy can be built using engineered expertise by adding specialised autonomic controllers to sub-elements of a system, but reaching level 5 autonomy will require whole-system approaches which can both integrate and then reason about an entire system end-to-end.

Summary. To summarise our study of research questions, the majority of contributions either seek to use autonomy to enhance performance, or alternatively explore research tools and ways to model and improve autonomy itself. The use of machine learning is the dominant strategy for autonomic controllers – perhaps unsurprising given the current popularity of machine learning in general in computer science – and most publications examine parametric adaptation in which parameters are automatically tuned.

As a potential source of future travel for the community, we observe that almost all contributions examine autonomy of level 4 or lower, with very few research works examining level 5. Intertwined with this, the vast majority of contributions focus their autonomy on a targeted sub-element of a system, generally requiring specialist expertise to identify and develop an autonomic controller for. From our perspective reaching level 5 autonomy requires a real shift here in focusing on ‘whole-system’ autonomy able to capture all elements of a system under a unified autonomic control paradigm.




Fig. 8. Results on the dominant methodologies used to answer research questions. N/A in this data is used for papers which do not have a clear methodology in a traditional sense, such as those reporting on exemplars of self-adaptive systems or on new platforms.




Fig. 9. The findings / results reported by the papers in our study set; N/A indicates that a paper does not have findings in a traditional sense, and is common for new platform contributions.

B. Methodology

We next examine the methodologies used to study the above research themes. This offers insight into the accepted methodological practice of the community, and the potential affects that this has on the robustness of reported findings. In general we avoid providing specific examples of research for each methodology in this section, as methodological practice is open to superficial value judgment.

The set of overall research methodologies used across the 210 papers in our study is shown in Fig. 8. Cumulatively over all three years of the study period, the majority of papers (78) used *real-world experiments* in which a real system is built and then subjected to realistic (or approximately realistic) conditions in which to study the effects of autonomy. The second most dominant methodology (68 papers) is *simulation*, in which a simulator is used to approximate both the system being studied and its operating conditions.

We view these results positively, in that real-world experiments tend to provide the most robust results (even though they can be harder to replicate). The volume of simulation-based studies still remains very high, however, indicating that many results are potentially unproven in more realistic settings.




Fig. 10. Data on how many papers have results that are easily repeatable by others, for example sharing source code and detailed experiment setup. The N/A category generally occurs for less traditional papers such as research proposals or conversations.

Fig. 9 shows the distribution of results / findings reported in the papers included in our study, coded into papers that report positive, neutral, or negative results. This potentially provides insight into the community’s view of negative versus positive results – and therefore how balanced the picture of research into autonomic computing is. It also potentially interacts with the above picture on methodological practice in terms of the detail resolution of a study.

The majority of contributions (125 papers) report entirely positive results, with little or no evidence of limitations or negative aspects of the approach or research question. Positive result publication bias is well-known to computer science in general, but it is still perhaps surprising that so little space is given to reporting on limitations.

The second highest cumulative category here is neutral results. We code the results as neutral if a paper reports mixed results and is not advocating a certain approach from the outset. A smaller number of papers then reports *mostly* positive results in which limitations are clearly identified, and just 2 papers report negative (or mostly negative) results.

We next study the repeatability of research results, with Fig. 10 showing our coded data into *yes*, *no*, or *if reimplemented*. In order to be judged as repeatable, the paper must either clearly explain its methodology (for systematic reviews) or be accompanied by working code and replication conditions (for contributions in which a built system is examined).

Our results show that the majority of eligible contributions (100 papers) in our study are *not* repeatable under these requirements; this includes papers with no source code, limited study setup detail, and also papers did feature a link to source code but that link was broken or otherwise unusable. A smaller proportion of papers (80) did have full repeatability, while a third group were judged to have sufficient methodological and pseudocode detail in the paper that they could have been replicated if re-implemented since no real code was provided.

In our view this particular result raises the most concern among all of our data: a community unable to easily replicate the majority of its reported results may struggle to grow and




Fig. 11. Data on the datasets used by the papers in our study set. N/A here indicates a paper for which a specific dataset is not used or relevant to the contribution.




Fig. 12. Our data on how many papers involved a real system being built or studied (note that we also count papers which take an existing real system and study it as ‘yes’). The N/A category here is typically more unusual papers for which a built system would not be appropriate or relevant.

break new ground in the long term. In future it may therefore be beneficial for the community to introduce new expected standards in the repeatability of submitted research in order to overcome this potential limitation. We note here that repeatability of results is, of course, a broad area of improvement needed within many communities in computer science and wider scientific fields, and we do not intend to single out the autonomic/self-adaptive community as particularly lacking here compared to other fields – but the data here is nonetheless compelling towards making positive changes in this field.

Looking at repeatability in more detail, Fig. 11 shows the kinds of datasets used in experimental studies across the set of contributions, where the nature of a contribution implies that a dataset is relevant (such as a cloud computing piece which uses a certain request pattern). The results here are generally positive, in that the majority of eligible papers cite a publicly-available dataset, or describe how to create a suitable synthetic dataset. The number of contributions for which a dataset is private or unavailable is still significant, however, creating its own impacts on repeatability.

Our final results, shown in Fig. 12-13, report how many contributions in our study feature (a) a real-system that was




Fig. 13. Our data on how many papers involved a comparison with a real system in their evaluation.

built; and (b) an empirical comparison of that system to an alternative real system. In the first of these results, the number of research contributions is almost evenly balanced between those that did (105) and did not (96) build a real system (we note there is some overlap here between methodologies that use real-world experimentation and simulation as in Fig. 8). The second result shows a different picture, demonstrating that relatively few contributions (43 papers) include a comparison with an alternative existing real system.

This indicates that the dominant methodology reports on studying an approach or system on its own terms or against a theoretical comparison, rather than comparing against a relevant alternative. While it is always interesting to observe a system in its own terms, and some types of research are notoriously difficult to meaningfully compare, it is often the case that a higher level of robustness in results and conclusions can be gained from attempting direct comparisons against other work. This situation may also be compounded by the repeatability results reported earlier, where relatively little research across these communities can be easily replicated – presumably making it difficult to then compare that work later.

Summary. The range of methodological practice in the field is varied, but significantly dominated by real-world experiments and simulation (split roughly equally, with slight bias towards real-world experiments). It is difficult to comment on the relative value of this split in itself, since more prototypical or far-future research may tend towards simulation and this is not necessarily a bad thing. We suggest that genuine warning signs in this data are (i) a significant bias towards purely positive results; (ii) the lack of repeatability of reported results; and (iii) the lack of direct comparison against other work.

C. Alt

We use this section to highlight research that falls outside of the dominant categories, as this level of detail is lost in our high-level graphs. We do this to help avoid normative reinforcement, which can be detrimental to the vibrancy and inclusiveness of a research community, and to demonstrate the potential diversity of the field and provoke broader thinking about self-adaptive systems research.

The selected works in this section are from the coding labeled ‘Alt’ in Fig. 2, and often appear in the coding ‘N/A’ in other graphs. We cover some of the broadest examples in these categories, featuring work that explores the nature of intelligence, the lessons we can learn from social systems, and wider technical challenge areas.

ElSaid, Desell, and Krutz present a discussion on whether adaptivity is a core property of intelligent systems [16], suggesting that the question is more complex than it may initially seem. The central question is distilled to the difficult topic of how intelligence itself is defined, and the authors define four kinds of intelligence in terms which are relevant to the presence of absence of adaptivity.

On a related theme, de Lemos and Grześ present a position on how self-adaptive behaviours relate to artificial intelligence [17] (rather than the popular trend of applying machine learning to self-adaptive systems). Here the authors propose that metrics of explainability and interpretability can be used as reward functions for self-adaptive logic to act on the internal models of machine learning such as neuron inhibition depending on observed drifts in data and apparent reward.

Continuing the theme of broader views on self-adaptive systems, the study of socio-technical, purely social, or democratic legal systems as self-adaptive or self-integrating examples is a long-running sub-area of the community. Barnes, Ekárt, and Lewis explore the idea of social action in socially situated agents [18], considering scenarios in which two (generally human) systems are introduced into the same world and how this impacts the goals and actions of those systems. This kind of study can lead to fresh insight on how technical systems can similarly negotiate their actions in unknown environments with other potentially competing actors.

In wider technical challenge areas, D’Angelo explores how self-adaptive theory can help to manage uncertainty in fully decentralised systems [19], exploring how full decentralisation impacts computation and how adaptive behaviours may aid with these impacts. Examining wider applicative theory, Sedgewick and de Lemos explore how blockchain technologies may be a powerful concept to facilitate distributed self-adaptive systems [20], particularly focusing on potential for enhanced resilience against cyber security attacks.

Summary. The research contributions outside of the dominant categories represent some of the most provocative and wide-ranging thinking in the papers included in our study. In our view these works represent some key centers for wider discussion and could potentially be given higher prominence by explicitly soliciting broader thinking in dedicated sessions.

D. Comparison to other data

In this section we compare our results to those of the closest prior work, which is the 2012 survey by Weyns et al. entitled ‘Claims and Supporting Evidence for Self-Adaptive Systems: A Literature Study’ [5]. The Weyns survey examines contributions to SEAMS between 2006 and 2011, and while the data is not directly comparable to ours due to the single- versus cross-community focus and slight differences

in research question, we make comparisons where applicable. We focus specifically on comparing application domains, autonomy objectives, methodology, and repeatability.

In application domains, our data shows that data centre applications dominate the contributions by a significant amount, followed by general SAS contributions considering generalised theory, then IoT, robotics, and a mixture of smaller topics. The Weyns data suggests a similar picture in 2012, with ‘service-based’ systems (roughly equivalent to data centre) being dominant by a significant margin, followed by robotics, embedded systems (cf. IoT), and then a range of smaller topics. The data here therefore matches quite closely in overall trends between the two time periods, and perhaps reinforces the question of what makes data centre systems so interesting in terms of autonomic computing research.

Regarding the objective of autonomy, our data shows performance (CPU) improvement as the dominant objective, followed by optimisation of the MAPE-K process, then distributed coordination and smaller topics. The Weyns data shows a slightly different picture, in which flexibility is the most common topic of study, followed by reliability, then performance and smaller topics. Besides this, the volume of contributions studying ways to model and reason about systems and autonomy is a notably dominant feature of both datasets – reinforcing our observation that the community is yet to find consensus on how to model systems for autonomy.

Examining methodology, the picture between our data and Weyns is very different. While we observe mostly real-world empirical methodologies, Weyns observes a very dominant example-driven methodology for evaluation, with very few contributions using either simulation or real-world study. This difference could be in part caused by the single-community focus of Weyns, revealing more common trends in that one community, but is also likely a reflection of maturation of the overall field towards a set of clear real-world exemplars for the study of autonomous systems.

Finally, on repeatability, our data resonates with that of Weyns in that the majority of contributions are not repeatable. Our data does show signs of improvement here, however: 38% of papers do have full repeatability in our study, compared to just 2% in the Weyns data. Again, making a direct comparison here is difficult due to variations in study methodology, but this potentially represents a positive trend in more recent research.

IV. DISCUSSION AND CONCLUSION

This paper asks the research questions: (i) what kinds of question are self-adaptive systems researchers trying to answer; (ii) what are the dominant methodologies used to answer these question; and (iii) how has the community progressed towards the vision of autonomic computing, guided by the questions that it asks and methodologies that it uses.

To answer these questions, we studied all papers published at the three leading venues for self-adaptive and autonomous systems (ICAC, SASO, and SEAMS) over the last three years.

In research questions, our findings indicate that there remains a significant focus on parametric adaptation, autonomic

controllers targeted at a specialised sub-element of a system, and on level 4 of autonomous control in which systems are still manually engineered and integrated to serve higher level business logic. Overall, these research characteristics result in a large number of papers which experiment with adding an autonomic control loop to a particular sub-system element to enhance a system according to a particular metric, and then duly reporting that the enhancement was successful.

While this form of research is very useful in providing more data on where autonomic control works well, it is worth reconsidering the grand challenge of reaching a higher level of autonomous system – in which entire new systems are automatically integrated from higher level logic and then automatically adapted to their environment. Achieving this appears to rely far more on structural adaptation and reasoning, rather than parametric, and requires autonomic control paradigms which are able to pervasively encompass an entire system rather than being designed for one specialised sub-element.

It is also interesting that there remains a large amount of research on *how to model* systems for autonomic or self-adaptive behaviours, how to model that autonomic behaviour itself, and on *platforms* which can support research. While it is, of course, perfectly healthy for a research community to strive for new tools and new theory enabling new ideas, we note that these areas put together represent more research than any other topic in the field. Finding standards is elusive at the best of times, but it may be useful for the community to further reflect on the tools already available to begin to unify research around leading platforms/exemplars.

On methodological practice, we find an almost even split between real-world experimentation and simulation as the top two approaches, with a theoretical analysis methodology being a strong third approach. The contributions in our study report a high number of positive results, with far fewer papers reporting neutral and negative results. In repeatability, the majority of publications do not supply enough detail or public implementation to support repeatability, and the majority of papers reporting on a real-world built system do not offer a direct comparison against an alternative system. Both of these factors provide cause for reflection on the potential robustness of conclusions, and the ability of researchers to collectively build on and compare their work over time.

Outside the dominant research categories we find a range of provocative contributions examining fundamental questions of intelligence and adaptation, and of social human systems and their relationship to adaptation in technical systems. In our view these works are of high value in broader thinking around the nature of autonomy, and if given a suitable platform are potential focal points for deeper discussion in the community.

As the world becomes ever more dominated by technology at ever larger scales, the ideas of autonomous and self-adaptive systems should be more relevant and necessary than ever. It is our hope that this study offers an opportunity for review and reflection on the current state of self-adaptive and autonomous systems research and its progress – a field we hope to see continue to grow in the most positive ways.

REFERENCES

- [1] J. O. Kephart and D. M. Chess, “The vision of autonomic computing,” *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [2] M. Salehie and L. Tahvildari, “Self-adaptive software: Landscape and research challenges,” *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 4, no. 2, p. 14, 2009.
- [3] B. H. C. Cheng, R. de Lemos, H. Giese, P. Inverardi, J. Magee, J. Andersson, B. Becker, N. Bencomo, Y. Brun, B. Cukic, G. Di Marzo Seruendo, S. Dustdar, A. Finkelstein, C. Gacek, K. Geihs, V. Grassi, G. Karsai, H. M. Kienle, J. Kramer, M. Litoiu, S. Malek, R. Mirandola, H. A. Müller, S. Park, M. Shaw, M. Tichy, M. Tivoli, D. Weyns, and J. Whittle, *Software Engineering for Self-Adaptive Systems: A Research Roadmap*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 1–26.
- [4] F. D. Macías-Escrivá, R. E. Haber, R. M. del Toro, and V. Hernández, “Self-adaptive systems: A survey of current approaches, research challenges and applications,” *Expert Syst. Appl.*, vol. 40, no. 18, pp. 7267–7279, 2013.
- [5] D. Weyns, M. U. Iftikhar, S. Malek, and J. Andersson, “Claims and supporting evidence for self-adaptive systems: A literature study,” in *2012 7th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*, 2012, pp. 89–98.
- [6] A. G. Ganek and T. A. Corbi, “The dawning of the autonomic computing era,” *IBM Systems Journal*, vol. 42, no. 1, pp. 5–18, 2003.
- [7] T. T. Tun, M. Yang, A. K. Bandara, Y. Yu, A. Nhlabatsi, N. Khan, K. M. Khan, and B. Nuseibeh, “Requirements and specifications for adaptive security: Concepts and analysis,” in *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 161–171.
- [8] K. Hakimzadeh and J. Dowling, “Ops-scale: Scalable and elastic cloud operations by a functional abstraction and feedback loops,” in *2019 IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, 2019, pp. 62–71.
- [9] F. Quin, D. Weyns, T. Bamelis, S. Singh Buttar, and S. Michiels, “Efficient analysis of large adaptation spaces in self-adaptive systems using machine learning,” in *2019 IEEE/ACM 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*, 2019, pp. 1–12.
- [10] B. Porter and R. Rodrigues Filho, “Distributed emergent software: Assembling, perceiving and learning systems at scale,” in *International Conference on Self-Adaptive and Self-Organizing Systems*. IEEE, June 2019.
- [11] R. Birke, J. F. Pérez, S. B. Mokhtar, N. Rameshan, and L. Y. Chen, “Chisel: Reshaping queries to trim latency in key-value stores,” in *2019 IEEE International Conference on Autonomic Computing (ICAC)*, 2019, pp. 42–51.
- [12] K. R. Jayaram, A. Gandhi, H. Xin, and S. Tao, “Adaptively accelerating map-reduce/spark with gpus: A case study,” in *2019 IEEE International Conference on Autonomic Computing (ICAC)*, 2019, pp. 105–114.
- [13] S. Huang, S. Liang, S. Fu, W. Shi, D. Tiwari, and H. Chen, “Characterizing disk health degradation and proactively protecting against disk failures for reliable storage systems,” in *2019 IEEE International Conference on Autonomic Computing (ICAC)*, 2019, pp. 157–166.
- [14] Y. M. Ramirez, V. Podolskiy, and M. Gerndt, “Capacity-driven scaling schedules derivation for coordinated elasticity of containers and virtual machines,” in *2019 IEEE International Conference on Autonomic Computing (ICAC)*, 2019, pp. 177–186.
- [15] R. Olacchia, J. Atlee, A. Legay, and U. Fahrenberg, “Trace checking for dynamic software product lines,” in *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 69–75.
- [16] A. ElSaid, T. Desell, and D. E. Krutz, “Is adaptivity a core property of intelligent systems? it depends,” in *Proceedings of the 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS ’19. IEEE Press, 2019, p. 153–154.
- [17] R. de Lemos and M. Grzeundefined, “Self-adaptive artificial intelligence,” in *Proceedings of the 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS ’19. IEEE Press, 2019, p. 155–156.

- [18] C. M. Barnes, A. Ekárt, and P. R. Lewis, "Social action in socially situated agents," in *2019 IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, 2019, pp. 97–106.
- [19] M. D'Angelo, "Decentralized self-adaptive computing at the edge," in *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 144–148.
- [20] P. E. Sedgewick and R. de Lemos, "Self-adaptation made easy with blockchains," in *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 192–193.