

Interview with David E. Smith

Pascal Bercher & Daniel Höller

KI - Künstliche Intelligenz

German Journal on Artificial Intelligence
- Organ des Fachbereichs "Künstliche
Intelligenz" der Gesellschaft für
Informatik e.V.

ISSN 0933-1875

Volume 30

Number 1

Künstl Intell (2016) 30:101-105

DOI 10.1007/s13218-015-0403-y



Your article is protected by copyright and all rights are held exclusively by Springer-Verlag Berlin Heidelberg. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Interview with David E. Smith

Pascal Bercher¹ · Daniel Höller¹

Published online: 27 October 2015
© Springer-Verlag Berlin Heidelberg 2015



David E. Smith is a senior Researcher in the Intelligent Systems Division at NASA Ames Research Center. He received his Ph.D. in 1985 from Stanford University, and spent time as a Research Associate at Stanford, a Scientist at the Rockwell Palo Alto Science Center, and a Visiting Scholar at the University of Washington

before joining NASA in 1997. Beginning in 1999, he served as the lead of the 18 member planning and scheduling group at NASA Ames for 6 years before abdicating to devote more time to research. Much of his research has focused on pushing the boundaries of AI planning technology to handle richer models of time, concurrency, exogenous events, uncertainty, and oversubscription.

Smith served as an Associate Editor for the Journal of Artificial Intelligence Research (JAIR) from 2001–2004, and as Guest Editor for the JAIR Special Issue and Special Track on the 3rd and 4th International Planning Competitions. He served on the JAIR Advisory Board 2004–2007. Smith was recognized as a AAAI Fellow in 2005, and served on the AAAI Executive Council 2007–2010.

✉ Pascal Bercher
pascal.bercher@uni-ulm.de

¹ Institute of Artificial Intelligence, Ulm University, Ulm, Germany

KI Would you like to tell us about your journey to NASA and your personal motivation for this?

While I was still at Stanford, I became increasingly interested in planning and scheduling. In particular, I thought there were opportunities to apply some of the ideas from my thesis work on controlling logical inference to the problem of controlling search in planning. I got the opportunity to work on planning when I came to the Rockwell Science Center's Palo Alto Laboratory (RPAL) in 1988. For the first several years at RPAL I participated in the DARPA planning initiative, which was largely focused on the development of planning and scheduling aids for logistics planning, and for planning non-combatant evacuation operations. At the time, there were a number of well known researchers in decision theory and uncertainty in AI at RPAL, and my collaboration with Mark Peot led to our foundational paper on conditional non-linear planning [17]. I was also involved in the development of algorithms and software for Design Sheet, a system developed by Ken Fertig to facilitate engineering conceptual design [7, 18]. My work on this was focused on the development and use of graph search algorithms to automatically partition sets of non-linear algebraic equations, and choose iteration variables for faster numeric solution and better convergence. Design Sheet was used by a number of Rockwell divisions, but is now owned and used by Boeing for aircraft design. In my final years at RPAL, I developed a system for newspaper imposition planning to help support the Goss printing press division of Rockwell. Newspaper presses are not monolithic—they consist of many different press units with different capabilities. The problem was to assign pages of a newspaper to the different press units, and determine the routing of the paper through the press. In essence, this was

a large constrained optimization problem where the objectives were to minimize press setup time and minimize the risk of paper breakage during a press run. The software underwent trials at several large newspapers including the Chicago Tribune, the Miami Herald, and Singapore Press Holdings. It was a product of Goss and Allen Bradley for several years.

In 1997 Brian Williams approached me about coming to work at NASA Ames. What ultimately convinced me was the compelling need for planning and scheduling technology for deep space missions and planetary exploration. For a distant spacecraft or rover, the communication delays are significant, and there is limited communication bandwidth. As a result, it is often impractical to make detailed decisions on Earth—the vehicle needs to be able to analyze the situation, and make timely decisions about what actions to perform. This is, I think, what motivates many who work in automated planning and scheduling at NASA.

KI *Can you tell us something about your current work and projects?*

Much of my recent work at NASA has been focused on aeronautics rather than space. While most people associate NASA with space exploration, the first A in NASA is for aeronautics. NASA does work in many areas of aeronautics, including aviation safety, human factors in the cockpit, air traffic management, aerodynamics, unmanned vehicles, and structures and materials. In part, my change in focus came about because I'm also an instrument-rated pilot, and have a personal interest in aviation safety and air traffic management. While commercial air travel is generally very safe, when an aircraft encounters challenging weather, is forced to divert, or faces an emergency, the workload in the cockpit is very high. The pilots must fly the aircraft, diagnose problems, and make decisions about what to do, which may involve choosing an alternate airport or landing site, and generating a good route to that site. Unfortunately, pilots often do not have time to consider all the pertinent options in these stressful situations, and as a result, they may not always make the best decisions. To help address this problem, I led a project focused on designing an *Emergency Landing Planner* (ELP) [15]. The system considers the aircraft situation (flight envelope), weather conditions, and airport characteristics, and produces a list of possible alternative airports/runways and routes to get to those options.

In 2010 we integrated the prototype software into the Flight Management System of a full motion simulator for twin engine transport aircraft, and did a study with five teams of professional airline pilots to evaluate the efficacy of the system [14]. The results and feedback were quite positive, but there are many technical and regulatory obstacles to providing such capability in a certified

commercial aircraft. More recently, we have focused on the use of this technology to assist in Single Pilot Operation (SPO). For SPO, a single pilot in the cockpit is assisted by a ground operator, who has the ability to provide dedicated assistance or control the aircraft when needed. This presents many challenges, including providing effective user interfaces, facilitating situational awareness between the pilot and ground operator, and providing helpful automation aids. For this purpose, the system had to be adapted to deal with many different aircraft, operating over a wide geographic area. In addition, we had to add capabilities to allow the operator to impose additional constraints, provide a broader range of optimization criteria, and allow more interactive route planning. These enhancements have been evaluated in recent pilot studies on SPO [8]. A smaller study has also been performed to evaluate the impact of the user interface and explanation capabilities on pilots' trust and reliance in the system [13]. As a result of this project, the problem of characterizing and evaluating trust and reliance for intelligent decision aids has become quite interesting to me.

While the above topics consume most of my time, I've also been involved in some smaller efforts with colleagues and graduate students. Principal among these is work on conflict avoidance for small Unmanned Aerial Vehicles [2], goal recognition [10], and planning under duration uncertainty [16].

KI *Did you or your colleagues participate in any NASA missions? How did those missions benefit from AI techniques? Was there AI planning involved?*

My work at NASA has been on the development of research prototypes to demonstrate new capabilities. I have therefore not directly participated in any NASA missions. However, several current and former members of the planning and scheduling group at NASA Ames have been directly involved in mission activities during my time at NASA. In particular, members of the group, in collaboration with the Jet Propulsion Laboratory (JPL), were responsible for the Mixed-Initiative Activity Planning Generator (MAPGEN) software used to do ground-based activity planning for the Mars Exploration Rovers (MER), Spirit and Opportunity. This software was designed to take daily experiment goals from the science team, along with operational and communications constraints (power, thermal, visibility, etc), and generate detailed activity plans for the rovers. MAPGEN was designed as a mixed-initiative planning system, where a human *Tactical Activity Planner* (TAP) would place and move around activities on timelines, and the system would show constraint violations, allowing the TAP to resolve them incrementally [4, 6]. MAPGEN permitted the TAP to "pin" activities in place, select any subset of the constraint violations, and ask the

planner to resolve those conflicts. When doing this, the planner used a “minimum perturbation” heuristic, which tried to minimize the changes to the plan, so that it would be easier for the user to track and understand those changes.

There are several reasons why MAPGEN’s planning was done interactively: (1) at the start of the planning process the science team might not specify or even be fully aware of all of their preferences, (2) the goal language was not rich enough to be able to express all those preferences, and (3) the plans are complex enough that large scale or global changes to the plan are difficult for the TAP and science team to understand. As a result, planning needed to be incremental, and each step needed to be of limited scope. A more detailed discussion of these issues can be found in [5, 21]. MAPGEN was considered highly successful and is still in use for the MER mission. Much of the technology has also been adopted to do science activity planning for more recent missions, including the Mars Phoenix Lander [11], the Mars Science Laboratory (Curiosity Rover) [1], and the Lunar Atmosphere and Dust Environment Explorer (LADEE) [3].

A second, and very different mission that the group has been heavily involved in, is the development of the Solar Array Constraint Engine (SACE) for optimization of the solar arrays on the International Space Station [19]. One might think that this just requires keeping the arrays pointed at the sun. However, it turns out to be much more complicated—there are many constraints that limit positioning and movement of the arrays, including motor constraints, shadowing, structural forces due to thermal expansion, plumes and water dumps. This is not a traditional task planning problem, but rather a kind of constrained optimization problem that involves choosing a sequence of operating modes (auto-tracking, tracking with one or more of the joints constrained, parking, or locking joints to specific angles) to guarantee that the sequence of operating modes remains safe with respect to the ISS power requirements, thermal, structural, and other constraints. The system was used as a tool by ISS flight controllers at Johnson Space Center from 2006 to 2014.

KI *How flexible are those systems, for instance with respect to changes of the current situation?*

The mixed-initiative design of MAPGEN and its successors has proven quite adaptable to different operational strategies across several different missions, as mentioned above. The possible activities of the rover or lander are encoded in a declarative fashion in an *Activity Dictionary*, where each activity is described by its conditions and effects, much as with PDDL 2.1 durative actions [12].

The SACE software solves a very specific targeted optimization problem. It is adaptable to changes in the

timing of different activities that impose constraints on the solar arrays. It is also adaptable to changing power needs of the ISS. However, it is specific to the mechanical operating characteristics and geometry of the current solar arrays on the ISS. Changing those mechanisms or geometry would require significant modification of the system. The principal motivation for developing the system was that additional solar arrays were added to the ISS in 2006, and the flight control team was no longer able to do the optimization for the expanded system using their existing methods and tools. Ultimately SACE was retired due to operating changes of the Russian thrusters that dramatically changed the constraints on the operating modes. Fixing this would have required substantial modification of the system.

KI *You mentioned the mixed-initiative system MAPGEN. How important is it in general to involve the user in decision processes of autonomous systems? Are there problems that need to be tackled?*

This is critical in many of the planning and scheduling applications I’ve seen. The reason is that the user may have hidden preferences or objectives that are not yet known or cannot be expressed to the system. In some cases, the whole planning process becomes iterative, with the users gradually refining their goals and preferences. There are several keys to making this work: (1) there needs to be good visualization and explanation of plans—users must be able to understand plans generated by the system so that they can recognize what they don’t like, and realize why; and (2) users need the ability to control and limit the planning process, perhaps by fixing or constraining certain parts of the plan, or specifying which parts of the problem should be revised by the planning system. The “pin” and “plan selected” mechanisms for MAPGEN are good but simple examples of the latter.

KI *What is more important for the deployed systems: their functionality or their usability? Do you have some ideas, how those systems could be improved with respect to their usability?*

Both functionality and usability were critical to the success of MAPGEN and SACE. For MAPGEN and its successors, the ability to describe activities and operational constraints and reason about them to detect constraint violations is critical functionality. However, the user interface was equally critical—the ability to effectively visualize and manipulate plans on timelines, and do planning in an interactive fashion was essential. For SACE the underlying optimization and evaluation of the sequence of operating modes is critical functionality. However, the ability to see that this plan respects the combined constraints, and to visualize the sequence and timing of mode

changes across many concurrent threads of activity was also critical. For my own work on the Emergency Landing Planner, the underlying route planning, route evaluation, and route optimization is critical functionality. However, displaying alternatives to the pilots and allowing them to investigate and manipulate those alternatives is equally critical. Our recent study on trust and reliance shows how the explanation of solutions influences pilot acceptance of the results [13].

KI *You mentioned the influence of explanations on the trust a user has in the system's results. What kind of questions can your system answer and how are these explanations generated? Can the developed techniques be generalized to answer questions about common AI plans?*

For a route planning system like the ELP, the primary thing the pilots want to ask is why a particular landing alternative or route is a good or bad option. There is a fairly complex risk model used by the system to generate and evaluate landing alternatives and routes. The explanations that we provide are based on (1) giving separate risk assessments for the en route portion of the flight, the approach, and the landing, and (2) identifying the factors that contribute most to the risk of each option. To illustrate, the kinds of explanations that have proven to be the most helpful are of the form: “the approach has a moderate chance of success because the (cloud) ceiling at the airport is close to the minimums for the approach.”

This is somewhat different than the explanations needed for task level planning, where a user might have questions about why certain actions are included in a plan, why actions are performed in a particular order, or why certain actions are constrained within a particular time window. In general, explanations for these questions require analyzing causal information about the structure of the plan. Even more difficult are hypothetical questions about why certain actions are not in a plan, or why an alternative plan will not work. For ELP's route planning, the causal structure of the plans is simple and obvious. What is not so obvious is the optimization criteria. Thus, our explanations are all related to why an alternative or route is good or bad. We also handle a limited form of hypothetical question—we can evaluate a plan generated or modified by the user, and provide an explanation like that above for the alternative. For task level planning, answering these kinds of questions can require plan validation or even additional planning to investigate the hypothetical alternative. Some of these issues are discussed in [21] for mission science activity planning.

KI *Companion Systems are characterized, among others, by being able to adapt to the individual user and its current situation. Is this also interesting for your team or NASA in*

general? Is there someone working on achieving these capabilities within the currently deployed systems?

Individualization has not received much attention in these systems. In some cases the user may be able to set simple preferences that govern default behavior, but there is no user modeling or automated adaptation. For the systems described above, providing additional capabilities or improving the overall user interface is often much higher priority. Part of the reason is that these systems are targeted at very specific technical domains and the users (operators, flight controllers, and pilots) are all professionals accustomed to dealing with technical software and decision aids. It is not clear how much individualization would benefit such highly trained users in these focused domains.

KI *You have won the ICAPS¹ Influential Paper Award for one of your papers and received an honorable mention for another. Would you like to tell us about these papers and their impact on current research in AI planning? Do you also think that those are your most interesting pieces of work? Would you like to mention some other interesting papers of yours?*

The first of those papers, “Conditional Nonlinear Planning” [17], considered the problem of how to do planning when actions have uncertain outcomes, but are fully observable. At the time, the prevailing approach to planning was “Nonlinear Planning,” which is now known as “Partial Order Planning”. We therefore had to expand this approach to deal with the inclusion of actions with uncertain outcomes. The key insight was that we needed to keep track of the *context* for the different branches in the plan, and the *purpose* for the different open subgoals, so that we could properly recognize and resolve conflicts. This paper sparked a lot of new research in the area of planning under uncertainty, which has now grown into a central topic at ICAPS conferences.

The second of those papers, “Choosing Objectives in Over-subscription Planning,” [20] came about as a result of work I was doing on planning activities for planetary rovers. The issue was that scientists are ambitious: they always give you more goals than can possibly be accomplished given the available time and resources. But some of those goals are more important than others. This is far different from the classical planning problem where there is a fixed goal that must be achieved. The paper introduced the notion of over-subscription planning (OP), where goals have utility, and the objective is to produce a plan that maximizes utility given resource bounds. This paper also sparked a lot of new work on Partial Satisfaction Planning (PSP) that considers goal utilities and action costs, but not

¹ ICAPS = International Conference for Automated Planning and Scheduling. <http://www.icaps-conference.org/>.

resource bounds. More recent work by Domshlak and Mirkis has shown that OP is quite different from PSP, and that different heuristics are required [9].

One additional paper that had considerable impact was the “Temporal Graphplan” paper [22]. This paper adapted the heuristic power of the Graphplan technique to the solution of planning problems where actions have duration. This sparked a good deal of interest and work on temporal planning, leading to the development of PDDL2.1 and the temporal track of the 2002 International Planning Competition.

All three of these papers explored new ground in the planning community by considering problems that fell outside of the bounds of what was being done at the time.

KI *What would you advise young scientists in the field of AI?*

There is now a vast array of powerful AI techniques for solving different kinds of problems. The trouble is that the assumptions behind an individual technique do not always match what is required for solving practical problems. My most interesting and useful insights have come from trying to apply existing techniques to real problems, and recognizing the limitations of those techniques. This leads to extending or generalizing those techniques, combining different techniques, or recognizing a new and different class of problems that require synthesizing new techniques altogether. Of course the danger of dealing with real problems is that you can get bogged down in the details of the application. My advice is to get involved with and really learn about an application, but constantly reflect on what the real problem is, and the assumptions and limitations of the techniques you are trying to use.

KI *Mr. Smith, thank you very much for this interesting interview!*

References

1. Aghevli A, Bachmann A, Bresina J, Greene K, Kanefsky B, Kurien J, McCurdy M, Morris P, Pyrzak G, Ratterman C, Vera A, Wragg S (2007) Planning applications for three Mars missions with ensemble. In: Proceedings of the 5th international workshop on planning and scheduling for space (IWPSS)
2. Barreiro J, Do M, Smith DE (2014) Intelligent UAS sense-and-avoid utilizing global constraints. In: Proceedings of the ICAPS-14 scheduling and planning applications workshop (SPARK)
3. Bresina J (2015) Activity planning for a lunar orbital mission. In: Proceedings of the 27th conference on innovative applications of artificial intelligence (IAAI). AAAI Press, pp 3887–3895
4. Bresina J, Jónsson A, Morris P, Rajan K (2005) Activity planning for the Mars exploration rovers. In: Proceedings of the 15th international conference on automated planning and scheduling (ICAPS). AAAI Press, pp 40–49
5. Bresina J, Jónsson A, Morris P, Rajan K (2005) Mixed-initiative planning in MAPGEN: capabilities and shortcomings. In: Proceedings of the ICAPS-05 workshop on mixed-initiative planning and scheduling
6. Bresina J, Morris P (2007) Mixed-initiative planning in space mission operations. *AI Mag* 28(2):75–88
7. Buckley M, Fertig K, Smith DE (1992) Design Sheet: An environment for facilitating flexible trade studies during conceptual design. In: Proceedings of the 1992 aerospace design conference. American Institute of Aeronautics and Astronautics (AIAA)
8. Dao AQ, Koltai K, Cals S, Brandt S, Lachter J, Matessa M, Smith DE, Battiste V, Johnson W (2015) Evaluation of a recommender system for single pilot operations. In: Proceedings of the 6th international conference on applied human factors and ergonomics (AHFE). Elsevier
9. Domshlak C, Mirkis V (2015) Deterministic oversubscription planning as heuristic search: abstractions and reformulations. *JAIR* 52:97–169
10. E-Martin Y, R-Moreno M, Smith DE (2015) A fast goal-recognition technique based on interaction estimates. In: Proceedings of the twenty-fourth international joint conference on artificial intelligence (IJCAI). AAAI Press, pp 761–768
11. Fox J, McCurdy M (2007) Activity planning for Phoenix Mars Lander mission. In: Proceedings of the IEEE aerospace conference. IEEE, pp 1–13
12. Fox M, Long D (2003) PDDL2.1: An extension to PDDL for expressing temporal planning domains. *JAIR* 20:61–124
13. Lyons J, Koltai K, Ho N, Johnson W, Smith DE, Shively RJ (2015) Engineering trust in complex automated systems. *Ergon Des* (in press)
14. Meuleau N, Neukom C, Plaunt C, Smith DE, Smith T (2011) The Emergency Landing Planner experiment. In: Proceedings of the ICAPS-11 scheduling and planning applications workshop (SPARK)
15. Meuleau N, Plaunt C, Smith DE, Smith T (2009) An emergency landing planner for damaged aircraft. In: Proceedings of the 21st innovative applications of artificial intelligence conference (IAAI). AAAI Press, pp 114–121
16. Micheli A, Do M, Smith DE (2015) Compiling away uncertainty in strong temporal planning with uncontrollable durations. In: Proceedings of the twenty-fourth international joint conference on artificial intelligence (IJCAI). AAAI Press, pp 1631–1637
17. Peot M, Smith DE (1992) Conditional nonlinear planning. In: Proceedings of the 1st international conference on artificial intelligence planning systems (AIPS). Morgan Kaufmann Publishers Inc., pp 189–197
18. Reddy S, Fertig K, Smith DE (1996) Constraint management methodology for conceptual design tradeoff studies. In: Proceedings of the 1996 ASME design engineering technical conferences and computers in engineering conference
19. Reddy S, Frank J, Iatauro M, Boyce M, Kürklü E, Ai-Chang M, Jónsson A (2011) Planning solar array operations on the International Space Station. *ACM Trans Intell Syst Technol (TIST)* 2(4):41
20. Smith DE (2004) Choosing objectives in over-subscription planning. In: Proceedings of the 14th international conference on automated planning and scheduling (ICAPS). AAAI Press, pp 393–401
21. Smith DE (2012) Planning as an iterative process. In: Proceedings of the 26th AAAI conference on artificial intelligence. AAAI Press, pp 2180–2185
22. Smith DE, Weld D (1999) Temporal planning with mutual exclusion reasoning. In: Proceedings of the 16th international joint conference on artificial intelligence (IJCAI). Morgan Kaufmann, pp 326–333