ICAPS `22 Industry and Application Track Panel

Chairs: **Mathijs de Weerdt** and Florent Teichteil-Koenigsbuch*

Panelists:

Christina Burt	Fortescue Metals Group
Hoong Chuin Lau	Singapore Management University
Cédric Pralet	ONERA
Andrea Micheli	Fondazione Bruno Kessler
Mark Boddy	Adventium Labs

**not present today*

Main goals

1. Getting ICAPS results to industry & society: Bringing state-of-the art knowledge to industry, governmental organisations
2. Supporting ICAPS researchers to work on relevant (/industry) problems

Program

1. Each panelist shares some experience in turn
    - Other panelist may react
    - Some will invite you (audience) to share your thoughts (via chat)
2. Discussing solutions
    - Please prepare your ideas and put these in the chat

These slides with additions from the discussion (and chat) will be made available afterwards via the website.

- *If you don't want your name included, please let me know (M.M.deWeerdt@tudelft.nl).*
- *If you want full name and affiliation (in case you're quoted), please include these in the chat.*

Panelists' thoughts

Christina Burt

- Developing software for mine-sites generating US$4-5 billion per quarter
- Our application is ripe for automated planning - real-time state-aware assignment of autonomous and manned haul trucks
- No individual wants to be responsible for production stopping because an algorithm failed on an edge case
- My experience with open source solvers (math programming and automated planning) makes me reluctant to use them in this setting
- **The lack of commercially developed solvers for automated planning makes the paradigm unappealing for our application**

Audience:

1. Do you have (other) examples of a critical planning applications in industry?
2. Do you have experience of using commercial solvers in your planning applications?

Please share in the chat!

1. Examples of a critical planning applications in industry

Christian Muise:

- Dialogue planning, e.g. Santos Teixeira, M., Dragoni, M. A Review of Plan-Based Approaches for Dialogue Management. Cogn Comput 14, 1019–1038 (2022)
- Business Process Management, e.g. Marrella, A., & Chakraborti, T. (2021). Applications of Automated Planning for Business Process Management. In International Conference on Business Process Management (pp. 30-36). Springer, Cham.

Mauro Vallati:

- Traffic control, e.g. McCluskey, T. L., Vallati, M., & Franco, S. (2017, August). Automated Planning for Urban Traffic Management. In IJCAI (pp. 5238-5240).

1. Using commercial solvers in your planning applications

Derek Long:

- At Schlumberger we have deployed a planner (built in house, based directly on academic software) that has been in the field now for a couple of years. It has drilled more than half-a-million feet under fully autonomous control (planning thousands of times) and we also have deployed "wireline" (electrical geological surveying) using the same tech.

Shirin Sohrabi:

- Example at IBM: https://www.ibm.com/cloud/automation/watson-orchestrate that uses AI Planning.

Jan Dolejsi, Schlumberger:

- We develop PDDL models and use internally developed temporal-numeric planner in Schlumberger. It is deployed on the edge and is automating jobs under human operator supervision. The robustness is always about two things: the solver and the model.

Sylvie Thiebaux:

- In a field trial, we have used both commercial and non-commercial optimisation solvers to coordinate a batteries owned by consumers. But this wasn't a 5B business!

Victor Paléologue:

- I used Fast Downward in a product btw, it was ok in terms of planning. However for other reasons this product did not come out.

Mark Boddy:

- As far as I can tell, the examples cited all involve planning experts implementing/deploying/maintaining applications, either from within a company or in close association. That is a pretty limited notion of a "commercially available" product.

Christian Muise: What makes a planner a "commercially developed solver"?

  - License? Robustness? Adoption? Slick UI? Documentation?

Victor Paléologue: How about companies attached to universities and research: can’t they offer some expertise to integrate your work?

SMU Singapore Management University | School of Computing and Information Systems

HC Lau

- Obstacles for success:
  1. First Mile Problem - Scientific community has limited access to domain knowhow and data
  2. Last Mile Problem - Pathway to deployment
- Successful domains for ICAPS: transportation and logistics, space and robotics, machine/manpower scheduling, path planning, etc
- “New” domains for planning and scheduling: supply chains, (trade) finance, advanced manufacturing, health care operations, etc

Audience:

3. What are obstacles for success in accessing industry domain know-how?

4. In what kind of industry has the collaboration been successful?

5. In which other “new” domains could Planning and Scheduling techniques be successful?

Please share in the chat!

3. Obstacles for success in accessing industry domain know-how

Derek Long (KCL, Schlumberger):

- I completely agree with this first and last mile characterization of the challenges - and my own experience is that this remains a challenge within a big company, even working from the inside and even with deployed successes to point to.
- AIPlan4EU is already second mile (not first mile)
- Generally, a first obstacle when working with domain expert is to overcome initial skepticism: you need to show you understand them, and show them the potential.

Jan Dolejsi, Schlumberger:

First mile: planutils + AIPlan4EU

http://aiplan4eu-project.eu/

Last mile: AIPlan4EU (if extended to plan execution) + commercially backed planner-as-a-service

Victor Paléologue:

- About deployment: this is a critical issue, and that would not have been possible if it was not open source, and I could not hack around the tools. For sure there is not a large offer of people capable of doing that.

Shirin Sohrabi:

- I think modeling is an important aspect. If input is incorrect, you cannot expect that the planner would do well.

Christian Muise:

- Biggest challenge we faced in building planning-based dialogue agent solution wasn't anything to do with the planner technology, but rather training people to think in a declarative way to specify the models (even if we don't show them PDDL). May be a problem restricted to settings where model solicitation is an element of the process, but it was our biggest hurdle by far.

Victor Paléologue:

+1 with Christian. People trained to produce such models are rare.

Derek Long (KCL, Schlumberger):

+1 to Christian - indeed modelling in general, declarative or procedural is not a skill that everyone has developed. And keeping a declarative model free from contamination with procedural heuristics is also a challenge...

Mark Boddy:

Such modeling is a skill that almost NOBODY has

developed. Our community and a few like it are very much exceptions.

Richard Freedman:

Some HCI and HRI researchers, especially within the AI-HRI subcommunity, have been looking into ways to train people to communicate in this way with their machines. I recall Sonja Chernova had a paper a few years ago that had some success training people to think about HTNs practically, for example.

Christian Muise:

Jan is putting together (with eventual help from the rest of us (time is a scarce resource ;))) some great educational modules to help onboard people for this.
https://github.com/AI-Planning/modeling-in-pddl

Mark Boddy:

There are also relevant modeling conventions and practices. Try asking a manufacturing manager whether a given piece of equipment should be modeled as a unary or metric resource…

4. Reasons for success are: easy access to data, underlying problems close to standard combinatorial optimization problems

Cedric Pralet

- **Name of this track** appropriate?

Novel application track (ICAPS 2013 -> 2021)	Industry & application track (ICAPS 2022)	Application track (several other AI conf.)

Audience: *please respond to the poll about the name of the track*

- **Personal experience in the "Space" domain** : successful app. for ICAPS because
  (1) automated planning & execution often inevitable
  (2) existence of a network of actors and events

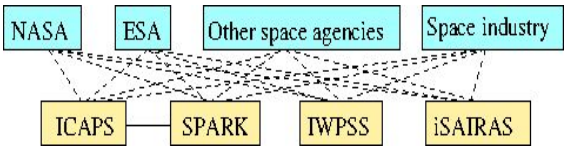

- Industrial end-users often prefer having an **adaptable** and **user-understandable method** for a **restricted problem** than a generic black-box planner
- Lessons can be learned from the **Operations Research** community (very successful in terms of collaborations with the industry)

All industrial and society applications are important, even the old but still difficult ones (the previous Application Tracks focused on *novel* applications).

Sylvie Thiebaux:

That is why the track is called industry and application track. We could call it industry OR application track?? :-)

=> Poll majority was to maintain "Industry & application track"

Derek Long (KCL, Schlumberger):

One slight problem in these tracks is that industry/application "war stories" are often short on the scientific contributions that are expected in a technical paper. This is partly because those parts are valuable and protected

by their owners and partly because the application does not necessarily involve solving problems of a type or in a form that would be interesting to present as a scientific contribution.
Alison Paredes:

+1 Derek. Especially the last part
Christian Muise:

Would a tailored call for that track help? Specifically calling out papers that "describe how you deployed planning tech in an industrial setting". Combined with a carefully selected (and informed) reviewer pool, this might push the needle.
Mark Boddy:

People have been pointing out that last part as a problem in the community for 30 years (1992, Earl Sacerdoti).

@Derek: apparently it's still a problem. So continuing to point it out is helpful!

@Christian: There is an innovative applications of AI conference (IAAI).
Alison Paredes:

Certainly in those 30 years someone has come up with some problems in the applications that ARE actually great settings for pushing fundamental planning research
Richard Freedman:

I think Christian is asking if we could make a IAAI within ICAPS?

Andrea Micheli

- Planning-based applications are **not stand-alone**
  - Part of a wider ICT system (e.g. digital-twin, robotic architecture, …)
  - The planning model is often not clean or clear-cut
    - (partially) constructed from data coming from other systems
    - requires adjustments, semantic attachments
    - sometimes, not static: model evolution
  - Planning alone is not enough: support execution!

- Planning tools are (often) not **designed for integration**
  - Command-line tools taking formal language specs vs library of functionalities
    - We need a pytorch/tensorflow of planning!
      - Encompassing major algorithms and tools
    - Embracing model engineering, evolution/updating
    - Visualization and debugging tools!

Audience:
Would you like planning-as-a-library or are you happy with command-line tools?
Please type 'library' xor 'command-line' in the chat.

Would you like planning-as-a-library or are you happy with command-line tools?

Jan Dolejsi, Schlumberger:

Planning as a library could be useful for prototyping. But eventually we need Planning-as-a-service to support the cloud, edge, or various architectures in between.

Victor Paléologue:

I want a library with an FFI, please (Foreign Function Interface)

Alison Paredes:

Or an API (Application Programming Interface)

Christian Muise:

library, and we kind of have it in various ways.

Jean Kiam:

For me, planning as a library

Victor Paléologue:

aka a portable library with a C interface

Christophe Guettier:

IAAI has a selection process with maturity levels, that could be of interest for the industry...

Sylvie Thiebaux:

It would be great if we had the tensorflow of planning and hopefully AIPlan4EU is going to build something like that!

Christian Muise:

But the UP (Unified Planning framework) offers something unique -- annotation of existing software to enable planning tech. Annotations on Python funcs, etc.

Andrea Micheli:

@Christian: not really sure what you mean by "annotate custom code to extract planning models". The UP is really a modeling+manipulation+solve library. Much like what you can do in OR: in addition to minizinc/SMT2 you have a solver API that allows you to interact with the solver without a printer+parser in the middle. The ambition of the UP is to make the interaction with a large range of solvers easy and convenient for application prototyping and development

Enrico Scala:

@Christian can you define better what you mean by integrating OO software within the UP? if you mean integrating externally defined procedures for effects or heuristic definitely yes

Christian Muise:

Yep, that and having existing python implementations start to generate PDDL elements (objects created corresponding to objects in the domain, etc). Happy to chat about it further offline, but really think it's an area UP can shine.

Victor Paléologue:

AIPlan4EU is not going to make a C library, I already asked them

Andrea Micheli:

@Victor: whay do you think a C-level API is so paramount? Wouldn't something like GRPC suffice at least for the desktop setting?

Christophe Guettier:

AIPlan4EU is going to integrate the AI4EU platform

that is part of the ICT49 cluster

Christian Muise:

Lapkt offers API access to piece things together (C++ / Python).

FD is moving in that direction afaik.

PaaS is currently under way (https://github.com/AI-Planning/planning-as-a-service),

AIPlan4EU really is filling the "annotate custom code to extract planning models" and effective execution (largely exploratory until now).

Derek Long (KCL, Schlumberger):

A challenge at the interface between academic/research software and industrial use is that the "boring stuff" of ensuring robustness, graceful failure, implementing all the bits that have no research value because others have already done it, are of very limited value to spend time on for an academic with shortage of time.

Shirin Sohrabi:

Michael Katz (IBM) has a number of softwares available as a docker image, worth noting.

https://researcher.watson.ibm.com/researcher/view_person_subpage.php?id=7989

Christian Muise:

I think the glue-that-binds would be a standard to integrate (e.g., "how do I call a remote server to get a plan"). AIPlan4EU / PaaS project (planutils-based) / IBM endpoints / etc. They don't need to all be the same deployed systems, but should hopefully "speak the same language".

Would be (very!) happy for Jan or AIPlan4EU to take charge on it ;)

Victor Paléologue:

Lapkt seems promising =)

Alison Paredes:

Re enterprise-wide platforms, maybe instead of targing traditional IT platform models like IBMs of the world, target niche groups within the company, that are capable of understanding planning's advantages

Christophe Guettier:

@Christian: while dealing with minizinc, have you tried minisearch? did it evolve?

Christian Muise:

@Andrea May have misunderstood some of the early examples, but if I have a custom set of software that defines things in an OO way, can I not integrated UP stuff to define the models pythonically? I thought this was on the roadmap for the work.

Victor Paléologue:

@Andrea, because we cannot necessarily rely on the network. And we should also be cautious as a business when we depend on cloud subscription or need for an extra device to make the product work (edge computing)

FFI is ugly, but it allows the library to be available to any language, including Python.

There's an in-between that tensorflow / pytorch found: you run everything in Python, but you can compile it into a native library that embeds your model for mobile device. This piece of library has a FFI to integrate with anything. I like this solution too.

Victor Paléologue:

Does AIPlan4EU plan to integrate with VSCode plug-in for PDDL? Or provide something competing with it?

Jan Dolejsi, Schlumberger, PDDL in VS Code:

VSCode is for the PDDL-first experience. It gives you support, while you model and helps you understand why the plans are the way they are. AIPlan4EU is for code(Python)-first. It can generate PDDL code, doesn't it? Then you could continue in PDDL in VS Code, if you find that

to be more compact representation of the model.

But AIPlan4EU provides the API for programmatic calls to planners, which should facilitate software integration and prototyping. And it proposes a binary encoding of domain, problem and plan, so the data on the wire is very small and the interaction can be much faster. And it works across platforms and across programming languages. You can prototype in Python and call solver in Java, deployed by planutils….

Alison Paredes:

@Jan you have described a consulting model. But you don't have to do all the consulting yourself. you can outsource it. No need to waste students time on it. You can outsource it with commercial certification.

Christian Muise:

To Alison's point, it would be amazing to see something from a certification-granting institution with a sufficient planning contingent in house. IBM comes to mind!

https://www.ibm.com/training/credentials

Derek Long (KCL, Schlumberger):

@AlisonP: Jan and I are inside the industry - no student time to spend! The commercial certification model only goes so far (we have run internal course several times - quite successfully, I think - but it is hard to sustain the development into actual use.

Alison Paredes:

@Derek Salesforce seems to be doing fine at that model

Victor Paléologue:

@Alison Paredes: I have a continual planning template for Pepper robots, running on Android

https://github.com/aldebaran/pddl-playground-for-pepper
Derek Long (KCL, Schlumberger):

@AlisonP: Yes - I need to understand better what they offer and how they make the commercial certification approach valuable (genuinely valuable - not simply a CV badge).

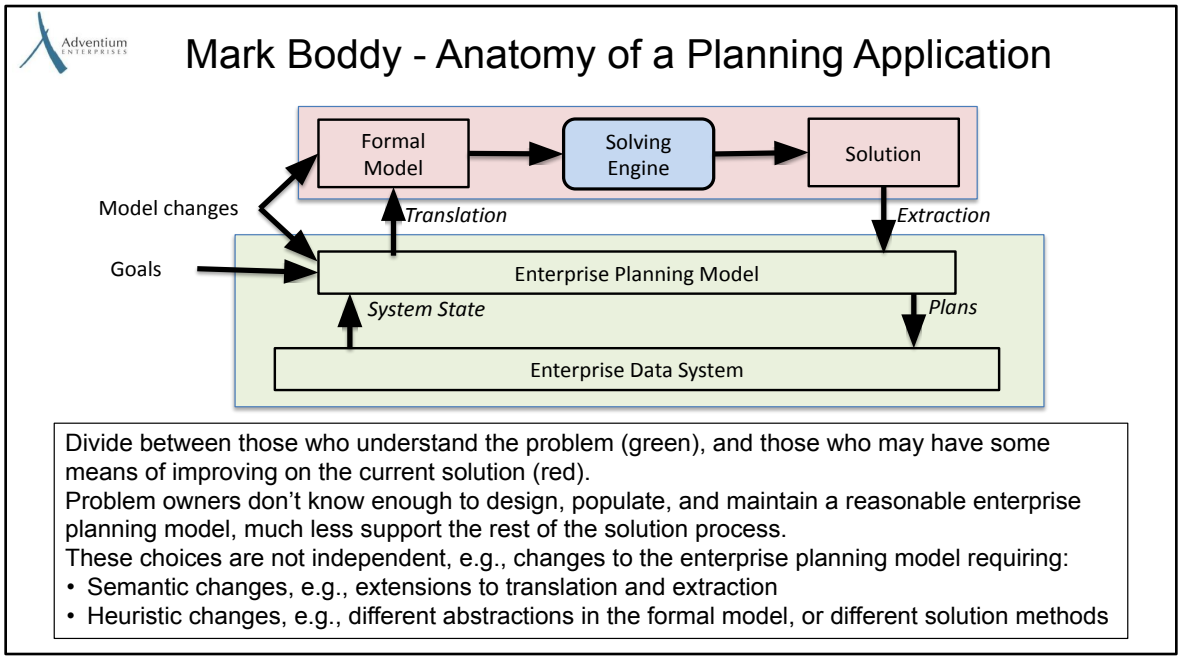

Red box is what (most) researchers mostly care about.

Green box is what the problem owners care about.

Problem owners don't know enough to design, populate, and maintain a reasonable enterprise planning model, much less the rest of the solution process.

These choices are not independent, e.g., changes to the enterprise planning model requiring:

- Semantic changes, e.g., extensions to translation and extraction
- Heuristic changes, e.g., different abstractions in the formal model, or different solution methods

- There is a divide between those who understand the problem,

- and those who have some understanding of possible means of improving on the current solution to that problem. There is no “planning toolbox” that can be used by a non-expert.
    - Need for proper process closely involving engineers and problem owners
    - Adapt the technology to the problem but not the opposite!
    - Input and output to stakeholders need to be well supported
    - Maintenance of the system is a challenge

Derek Long (KCL, Schlumberger):

Hector is right - you need to understand the fundamentals and the "right" way to tackle a problem is from the simplest form. But, scientists gonna science... The industrial applications have their own tempo.

Christian Muise:

"scientists gonna science" 😄

@Derek Hope to change that with an upcoming (grad) course on planning -- it puts it in a position to be useful for all grad students, and not just those doing planning. It'll be a mix of theory, approaches, and models for different formalisms.

We're starting to collect things here: http://education.planning.domains/

...and I'd highlight Jan's awesome start to modeling: https://github.com/AI-Planning/modeling-in-pddl

LAU Hoong Chuin:

The onus is on the instructors to make the course more applied (set assignments that are closer to industry) so that

students will have more hands-on opportunities to practise modeling in real settings.

Alison Paredes:

As I mentioned before, a certificate and consulting model ike Salesforce does (assuming there is some profit behind it), is something ya'll should coonisder

Mark Boddy:

+1 to Alison.

Alison Paredes:

Salesforce does an amazing job of getting totally untechnical people to do object-oriented modeling

Mauro Vallati:

Tools and approaches for supporting the overall knowledge engineering process -- this is what KEPS is about, in fact

Mark Boddy:

KEPS is very relevant and useful. But are the tools used (or useful) only for planning experts?

Mauro Vallati:

@Mark, good point. the focus is really both experts and non experts

Erez Karpas:

There could be many ways to model the same problem in PDDL, where one could be solved efficiently by state-of-the-art planners, and others could take forever. To know which model would be solved efficiently requires some understanding of the planning algorithms being used (even ignoring the choice of

which planning algorithm to use)
Enrico Scala:
+1 to Erez
Andrea Micheli:
@Erez: yes. And we need to make it easy to try different planning engines easily without recoding/hacking
Mark Boddy:
@Erez: Yes, absolutely. That's why you need a team that spans the range of expertise.
People used to think expert systems meant you didn't have to know how the underlying inference worked. Wrong then, wrong now.
Cédric Pralet:
In the OR community, some papers compare several MILP models for the same problem and analyze which model is the best. It would be interesting to provide such insights for planning models.
Mark Boddy:
@Cedric (sorry for misspelling): would such a paper be accepted at ICAPS? Should it be?
Cédric Pralet:
@Mark: I do not have the answer to the first question. For the second one, yes, it is valuable work I think.
Christophe Guettier:
@Cedric: that is the concept of a marketplace...
Enrico Scala:
this net distinction between models (only physics) and algorithms doesn't always work.
Derek Long (KCL, Schlumberger):
+1 Cederic - I think it is interesting that in the MILP/CP

community, models are considered a relevant research contribution, but in the planning community we are (partly intentionally - we want to claim that our tools are generic) less interested in the modelling.

Andrea Micheli:

Interesting point. I see the planning model as a piece of code that should evolve using practices much like in standard software engineering. So optimizations/profiling is just a part of the development process. Would you agree?

Mark Boddy:

@Andrea: Not completely agree, no. As discussed in this chat, there is an interplay between solution algorithm/method/tools and the model. There is no one "right" model for a given application.

Performance can be improved by focus on model, heuristics, *or* algorithm.

Mauro Vallati:

There is no right model, but there are definitely many wrong models -- KE should primarily help you to avoid those

Mauro Vallati:

algorithms and models will then co-evolve

Erez Karpas:

+1 Mauro

Mark Boddy:

@Mauro: KEPS is very relevant and useful. But are the tools used (or useful) only for planning experts?

Moving towards solutions

1. Supporting non-planning experts
2. Supporting ICAPS researchers

Supporting non-planning experts

- Develop professional software and support for solvers (Christina, Andrea, ...)
  - a library of functionalities (much like tensorflow or pytorch for neural-networks or e.g. ray rllib for RL) (Andrea)
  - “Plan util” tools for modeling/engineering, plan manipulation, visualization, validation (Andrea, Mark, Jan) -> ecosystem
- Build simple examples of continual planning & execution solutions for basic problems like dynamic OR problems (Cédric)
- Make a team of planning/execution experts and domain experts or forge long-term collaborations (like in space, Cedric, Mark)

Audience:

1. Do you have other ideas to support *non-planning* experts?

Please share in the chat!

Supporting (and incentivizing) ICAPS researchers

- Accumulate realistic benchmarks (add to IPC, SPARK):
  - Easy access to data (HC)
  - Formulate (the core of) industry problems in a more abstract and formal way (HC)
  - Consider problems as they are, instead of making them fit a given formalism (Florent)
  - Introduce “New” domains: supply chains (risk and resilience), trade finance, advanced manufacturing and health care operations (HC)
- Take lessons from OR (Cedric):
  - industrial competitions,
  - high acceptance rates for events gathering many academic and industrial participants,
  - include research topics including a holistic view of planning in practice

Audience:

2. Do you have other ideas to support ICAPS researchers?

Please share in the chat!

Christophe Guettier:

Realistic benchmarks are really hard to share, generally strategic for a business model

and business model <=> problem model

15:59:04 From Andrea Micheli:

@Mark: I agree, but if you see the application as a whole, I think that the model and the solver are both parts of it. So, it should be possible to make them co-evolve safely provided we follow a good development practice (CI + tests). Do we need dedicated "PlanningOps"?

Jan Dolejsi, Schlumberger:

Yes to PlanningOps. And hoping AIPlan4EU will evolve to that direction.

Relationship between industry and research: how to work

- Need for proper process closely involving engineers and problem owners:
    a. Understanding the application
    b. Describing the problem to be validated with the stakeholder
    c. Formulating that description in a representation, as input for an *appropriate* solution method
    d. Extraction of results from that solution method and presentation in a format which the stakeholder can understand and is happy with.
- Adapt the technology to the problem but not the opposite!
- Researchers should address continued improvement for well-known problems.
- Improve stakeholders' view of planning and scheduling:
    a. Focus on applications, not solution methods.
    b. Highlight success stories, especially those with major impact.

Presented by Mark

"Make ICAPS research great (again)" - seen by outsiders (Florent)

Alison Paredes:

Rebrand planning, just call it AI (like ML does)

Jan Dolejsi, Schlumberger:

Machine Reasoning

Derek Long (KCL, Schlumberger):

@AlisonP: My experience is that if you say you work on AI, it is immediately assumed that you work on ML...

Jakub Med:

+1 Derek

Mauro Vallati:

+1 Derek

Mauro Vallati:

I tried saying that planning is "model-based AI", but this is not as sexy as ML

Richard Freedman:

The issue is that ML researchers have convinced the media that their work is the entirety of AI. Rebranding to fight that will not help. I think it is a challenge of raising awareness and better communicating our science to the public.

Alison Paredes:

Planning is just ML when you don't have enough data ;-)

Christian Muise:

We need a marketing team in the community...

Mauro Vallati:

not only branding, also kind of difference in terms of fire power?

Derek Long (KCL, Schlumberger):

Wow - this has been a very active session - lots of interest here. Seems that it has hit a nerve in the community - and lots of participants in the audience too. I guess there is a strong sense of frustration that the work we do has had too little success in getting adopted and used outside our community.