

Timetabling Research: A Progress Report

Jeffrey H. Kingston

The University of Sydney

Introduction

The first PATAT conference was over 25 years ago, in 1995.

- How have the sub-disciplines changed since then?
- What about the solvers?
- Has our insight into timetabling deepened?
- Where do we go from here?

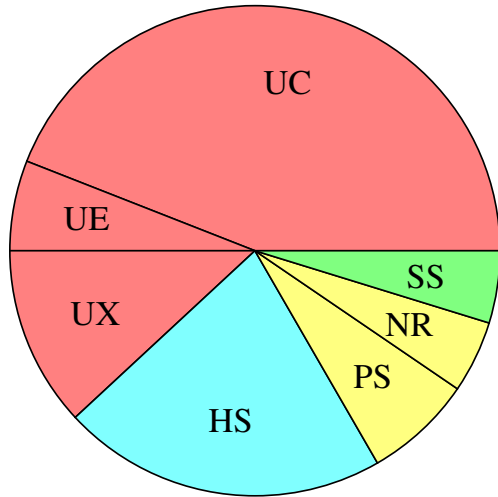
Methodology

Take 3 pairs of PATAT conferences and classify their papers:


- 1995 and 1997
- 2006 and 2008
- 2016 and 2018

All papers included, including plenaries, system demonstrations, etc.

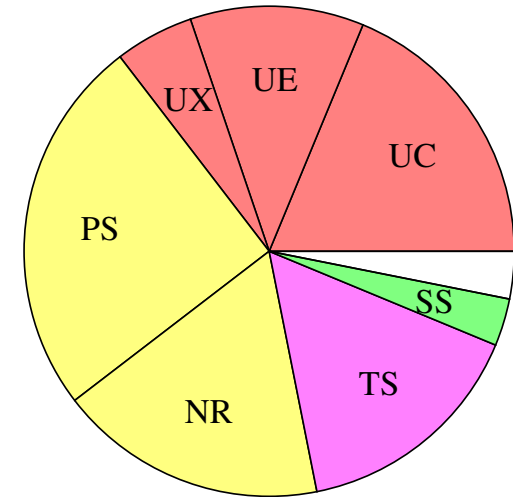
Papers from each sub-discipline



1995 and 1997

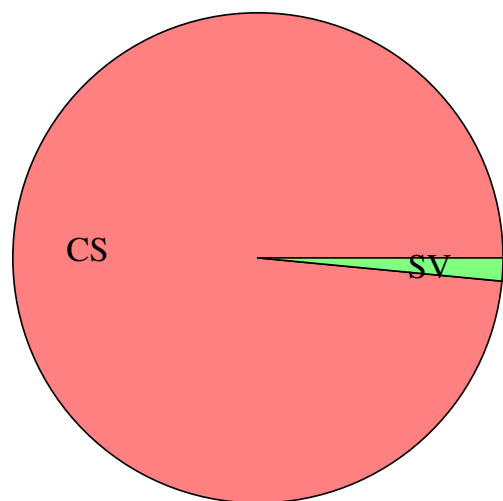


2006 and 2008




2016 and 2018

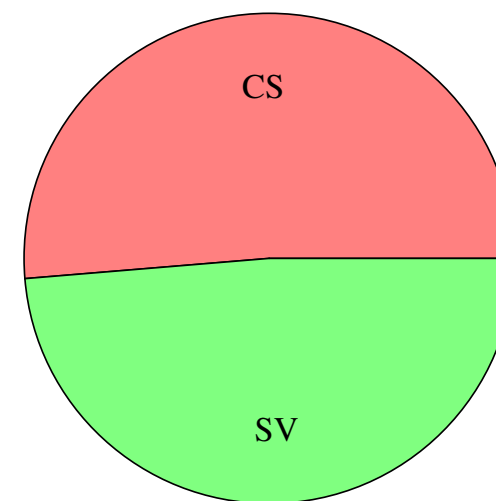
Case study papers vs solver papers



1995 and 1997




2006 and 2008




2016 and 2018

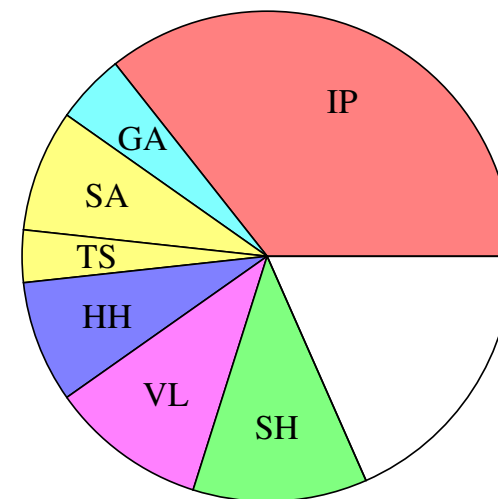
Solver types



1995 and 1997



2006 and 2008



2016 and 2018

Progress within sub-disciplines

Stage 1

Just a few case study papers; scope unclear

Stage 2

Plenty of case study papers; scope becoming clear

Stage 3

Standard data sets, competitions, and solver papers; scope clear

Stage 4

Decline; fewer papers; no clear research agenda

What constitutes progress depends on the sub-discipline's stage.

Example: personnel scheduling excluding nurse rostering

- Many papers (see figure)
- But no general picture, no standard data sets
- Conclusion: Stage 2 but ready for Stage 3

Discussions of the other sub-disciplines in the paper.

Insight into the timetabling problem

Solutions are better, but do we understand timetabling better?

- Solving: blocked by NP-completeness
- Specification: steady improvement
- Insightful papers: few and scattered

Insight has deepened, but only very slowly.

Moving forward

Our goal

Automated timetabling seeks to help people find high-quality timetables quickly and reliably wherever they are needed.


Case study papers can become backward-looking

When the scope of a sub-discipline is already clear

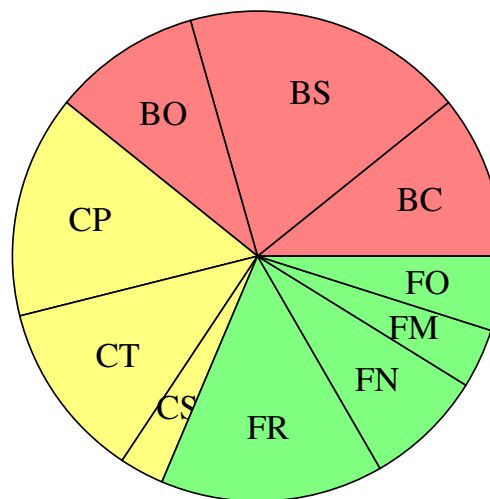
Solver papers can become backward-looking

When good, real-world data sets are already being solved to near optimality.

Diminishing returns example: instance BGHS98



Forward-looking and backward-looking papers



2016 and 2018

Suggestions for forward-looking papers

- Large case studies
- Faster and more robust solvers
- Minimal perturbation problems
- Infrastructure (data formats, data sets, competitions, ...)
- Dissemination of timetabling expertise

Overall theme: recommit to practice.

Appendix: success in practice

Academia is biased against practice. We need a precise, challenging definition:

*A solver is **successful in practice** if, on every instance that is likely to be encountered in practice, it finds a solution whose cost is within 10% of the best known when run for 5 minutes, and within 5% of the best known when run for 60 minutes.*

The challenge is spread across the three criteria for success in practice: good solution quality, moderate running time, and robustness.