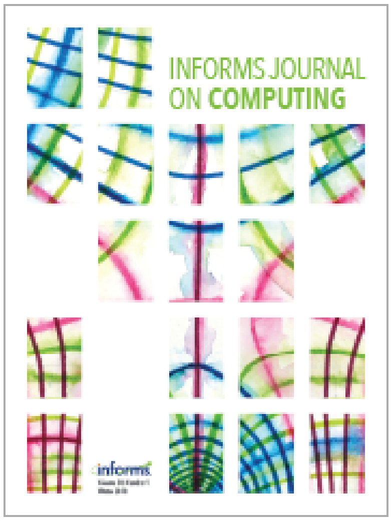

INFORMS Journal on Computing

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Unforeseen Consequences of "Tabu" Choices—A Retrospective

Fred W. Glover

INFORMS JOURNAL ON COMPUTING
Vol. 34, No. 3, May–June 2022, pp. 1306–1308
ISSN 1091-9856 (print), ISSN 1526-5528 (online)
http://pubsonline.informs.org/journal/ijoc

Unforeseen Consequences of "Tabu" Choices—A Retrospective

Fred W. Glover[a]

[a] Entanglement, Inc., Boulder, Colorado 80302

Contact: fred@entanglement.ai, https://orcid.org/0000-0001-6945-0438 (FWG)

Published Online in Articles in Advance:
April 5, 2022

https://doi.org/10.1287/ijoc.2022.1187

This retrospective is written upon receiving The *IJOC* Test of Time Award for the papers "Tabu Search—Part 1 and —Part 2," a delightful capstone to an unanticipated (and exceptionally fun) adventure.

I am prompted to tell you a story about the circumstances surrounding this adventure that is entirely personal and has no academic merit. In prelude, I confess that I have tried several times to get tabu search to be called adaptive memory programming (AMP) but entirely without success. So, I'll stick with tabu search (and its acronym TS). To tell the story underlying the method, I'll have to step back in time.

The year was 1970, and I had just joined the Business School at the University of Colorado, Boulder. A colleague who was instrumental in arranging this new position informed me of a combinatorial optimization problem the Bell System wanted to solve, with too many binary variables to be approachable by available methods. A procedure to obtain a local optimum was apparent, but Bell was casting about to discover a useful way to go farther.

Reflecting on this challenge brought back the memory of a psychology experiment I'd done a few years earlier as a doctoral student in an artificial intelligence (AI) class. (AI was enjoying a surge of popularity at the time, and its seesaw trajectory of dropping into the doldrums in the mid-1970s before dramatically recovering in the 1980s lay still in the future.) The purpose of the experiment was to discover how graduate students with varied backgrounds might go about finding good solutions to a combinatorial optimization problem I disguised as a word problem.

The outcome was surprising. Independent of their backgrounds, the students all followed a similar scheme in the way they generated a series of choices for probing various solution possibilities. There was a difference, however, between the procedures used by students who were more and less successful at finding the best solutions. The less successful students would either forget or discount the impact of choices made at an earlier stage and would effectively become "locked into" a set of decisions that proved less than desirable. The successful students initially made choices like those of their unsuccessful counterparts but more readily identified and modified previous choices that were relevant to change, ultimately finding their way to better solutions.

Two features of this experiment struck me as particularly intriguing. First, it offered an alternative to the Freudian view that people become stuck in unproductive behavior primarily because of repressing earlier experiences. By contrast, the experiment suggested that access to better decisions can just as easily be lost by simply forgetting or discounting the relevance of earlier choices. Second, these students (who arguably were all decent problem solvers) proceeded by temporarily enforcing their choices and then gradually revising them over time. The students constructed their solution paths more flexibly than by generating branches in a tree, and the difference between more and less successful approaches lay in an ability to recover and change previous decisions that were influential.

Accordingly, I decided to see if I could devise a scheme to tackle the Bell System problems that took advantage of the second feature of the experiment, temporarily enforcing decisions that appeared attractive and then, after a number of periods (a parameter), releasing the decisions as a basis for going further. (The approach was accompanied by a simple aspiration criterion that allowed decisions to be reversed before their "enforcement period" expired if this reversal produced a solution better than the previous best.) Because I was using a computer, I didn't have to worry about the possibility of forgetting earlier decisions that were appropriate candidates to revise.

On an impulse, referring to the title of a song called "Tabu" that was popular when I was young, I called the restrictions that prevented decisions from being changed "tabu restrictions." However, I didn't assign a name to the method itself.

This primitive form of TS worked surprisingly well, getting significantly better results for the Bell problems than the best previously discovered. However, my associates and I were caught off guard when the Bell System group dismissed these results as not possibly being as good as claimed. We were politely informed that Bell researchers were much better versed in ways to solve the problem than we were. The subtext was that academics should stick to academic pursuits and leave practical applications to those better qualified for dealing with them.

Although this response was disappointing, research interests that claimed a higher priority propelled me forward without much thought about the Bell rejection. Throughout this period, I was deeply involved in integer programming (IP) and network optimization, and I speculated that the approach devised for the Bell System application, which was a special instance of an IP problem, might be applicable to other integer programming problems. Drawing on ideas I'd been contemplating for IP heuristics using surrogate constraints, I devised supplementary strategies for handling more general applications. The resulting procedures included a forerunner of the TS path relinking strategy called scatter search, an intensification procedure based on a notion of strongly determined and consistent variables, and the first embodiment of the TS strategic oscillation approach. These elements were assembled in a short tutorial paper featuring the strategic oscillation approach, which was illustrated using separate tabus for alternating phases of oscillation, without using tabu terminology. Published in *Decision Sciences* in 1977, the paper went largely unnoticed.[1]

Meanwhile, network optimization was gaining increased attention in optimization circles, and I became especially intrigued by the discovery that more general models called *netforms* (for network-related formulations) embraced a wide variety of practical combinatorial optimization applications outside the pure network domain. Appreciating this development and the importance of network optimization generally, I cofounded a company for pursuing these applications called ARC (for analysis, research, and computation). My fledgling TS ideas found occasion to be applied in this setting when ARC was contacted by an architectural firm with a network-related problem arising in architectural design and space planning. A study describing this problem and the method for solving it was published in the *Annals of Operations Research* in 1985, becoming the first paper to introduce the "tabu search" name and terminology. As I was now becoming accustomed to expect, this paper too went unnoticed.[2]

This chain of events that kept the TS ideas from view was finally broken by informal communications with colleagues who picked up on the ideas and conducted experiments with them. The following year, things began to change. For an article invited by *Computers and Operations Research* titled "Future Paths for Integer Programming and Links to Artificial Intelligence," I included a brief section on tabu search that reported these experiments. In contrast with my previous forays into this area, the article attracted a wide readership, and the tabu search name and terminology no longer remained in obscurity.[3]

This somewhat erratic progression of events set the stage for the Tabu Search—Part 1 and —Part 2 papers. Although the "Future Paths" paper was well received, many of the key ideas underlying TS (including those of the earlier 1977 paper) were not yet assembled in a single place. In 1988, the *INFORMS Journal on Computing* (then called the *ORSA Journal on Computing*) presented an opportunity to remedy this with an invitation to write a paper devoted solely to TS. The resulting two-part publication inaugurated an interest in TS beyond all expectation, launching a series of research papers that transformed the tabu search method into an object of serious study.

It is impossible to begin to itemize the contributions by leading researchers that have enriched the TS literature following the release of the Part 1 and Part 2 *IJOC* papers. A rough idea can perhaps be gained from the Reference section of the 1997 *Tabu Search* book coauthored with Manuel Laguna and from more recent references that can be accessed by Googling "tabu search." (Some selectivity will be prudent, as the number of TS results returned by Google typically exceeds a million. A smaller selection can be found in the link https://tinyurl.com/TS-Pubs.)

You may have noticed there is sometimes a gap between the introduction of ideas and their implementation. This is particularly true in the case of tabu search, which has the upside that the advances made available through TS still have no end in sight. Although a quarter of a century has passed since the publication of the TS, Part 1 and Part 2, papers and the elaboration of their ideas in the *Tabu Search* book, a variety of strategies proposed in these references have yet to be explored. Additional TS-based strategies proposed in more recent publications likewise remain to be tested. It is tempting to speculate that a collection of these "overlooked strategies," judiciously assembled, could prove useful, both as an impetus to empirical research and as a foundation for designing other new strategies for challenging combinatorial optimization problems.[4]

Endnotes

[1] When this paper was written, heuristics were accorded a significantly lower status than "algorithms" (then the term for finitely

convergent methods), and I felt compelled to challenge this view with the preamble:

> Heuristic solution methods for integer programming have maintained a noticeably separate existence from algorithms. Algorithms have long constituted the more respectable side of the family, assuring an optimal solution in a finite number of steps. Methods that merely claim to be clever, and do not boast an entourage of supporting theorems and proofs, are accorded a lower status. Algorithms are conceived in analytic purity in the high citadels of academic research, heuristics are midwifed by expediency in the dark corners of the practitioner's lair.
>
> Recently, however, there has been a growing recognition that the algorithms are not always successful, and that their heuristic cousins deserve a chance to prove their mettle. Partly this comes from an emerging awareness that algorithms and heuristics are not as different as once supposed—algorithms, after all, are merely fastidious heuristics in which epsilons and deltas abide by the dictates of mathematical etiquette. It may even be said that algorithms exhibit a somewhat compulsive aspect, being denied the freedom that would allow an occasional inconsistency or an exception to ultimate convergence. (Unfortunately, ultimate convergence sometimes acquires a religious significance; it seems not to happen in this world.)
>
> The heuristic approach, robust and boisterous, may have special advantages in terrain too rugged or varied for algorithms. In fact, those who are fond of blurring distinctions suggest that an algorithm worth its salt is one with "heuristic power."

I was not very effective in enlisting supporters for this view, which impelled me some years later to coin the term "metaheuristics" to refer to high-level heuristics that can encompass and guide other heuristics. The metaheuristics term has fortunately been more effective in gaining support.

[2] The main problem the architectural design company wanted to solve was a mixed integer goal problem that may be viewed as a special type of constrained clustering problem, which today might be classified as an AI type of problem. Now that we have more advanced TS procedures perhaps this problem might usefully be revisited for its relevance to the AI setting.

[3] This paper did not hesitate to wax philosophical. Before getting to the tabu search topic, it made critical observations about the simulated annealing approach that would undoubtedly have raised the hackles of any Simulated Annealing (SA) proponents who might have chanced to read it. Ironically, the suggestion about how simulated annealing might be improved was a strategy that several SA researchers began to incorporate in their methods several years later.

[4] For example, the final chapter of the TS book illustrated a special clustering procedure as a way to improve TS methods, and other forms of clustering and biclustering have been similarly proposed for this purpose in following years. This is one of those instances where testing to verify or negate the merits of such proposals has fallen behind and that would seem a fertile avenue for exploration. Proposals about diversification procedures constitute another area where this "testing gap" conspicuously appears.

References

Glover F (1977) Heuristics for integer programming using surrogate constraints. *Decision Sci.* 8(1):156–166.

Glover F (1986) Future paths for integer programming and links to artificial intelligence. *Comput. Oper. Res.* 13(5):533–549.

Glover F (1989) Tabu search—Part I. *ORSA J. Comput.* 1(3):190–206.

Glover F (1990) Tabu search—Part II. *ORSA J. Comput.* 2(1):4–32.

Glover F, Laguna M (1997) *Tabu Search* (Kluwer Academic Publishers, Springer).

Glover F, McMillan C, Novick B (1985) Interactive decision software and computer graphics for architectural and space planning. *Ann. Oper. Res.* 5:557–573.