ELSEVIER

European Journal of Operational Research 106 (1998) 221–225

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

Tabu Search – Wellsprings and Challenges

Fred Glover [1]

School of Business, University of Colorado, CB 419, Boulder, CO 80309-0419, USA

1. Introduction

It is inevitable when discussing applications of tabu search in optimization that someone will ask how the name of the approach came into being. Some have imagined it to embody an acronym (as for "Trajectory Analysis via Biocybernetic Unification" or "Totally Awesome and Basically Unbelievable") while others have imagined it to carry a slightly risqué connotation (as where "tabu search" might represent the name of a new party game). In truth the term has less pretentious (and perhaps less exciting) origins, which result from a discovery that came about some years ago.

As part of a project for a graduate artificial intelligence course in the early 1960s, my classmates and I were challenged to devise a computer model (for the archaic computers of those days) to emulate some aspect of problem solving. I had been doing research on integer programming at the time, and consequently was curious to discover how my friends might solve an integer programming problem – under conditions where they were not told the nature of formal solution approaches, but were allowed to proceed solely on the basis of their own ingenuity. Accordingly, I sought to pose the problem in the form of a puzzle, by devising examples expressed in words rather than in mathematical notation (e.g., given specific columns of numbers, each with an attached profit value, select bounded integer multiples to create a new column whose entries do not exceed specified quantities, while seeking to maximize the associated total profit). Then I asked my friends to talk aloud as they tried to come up with their best answers, and to retain any scribbles on paper that they made in the process, as a basis for translating their behavior into a computer model.

As it happened, there were intriguing similarities in the procedures my friends used to find the best solutions. Uniformly, they adopted an approach that "vaulted past" the kinds of steps usually postulated to be the basis of logical thinking. Yet their approach was far from being unsystematic. They all erected certain patterns that effectively avoided reintroducing elements of constructions recently tried, except where this produced an outcome that appeared too good to pass up. A few of my friends played this game very well, and improved their strategy by isolating attractive possibilities along the way, which they revisited when a particular vein of exploration became unpromising. After extended intervals where no improvements materialized, these efforts were sometimes amplified to bring about more radical changes – changes that conspicuously were not random, but that were structured to create responses of greater "influence" relative to criteria that grew out of obstacles encountered.

The manner in which my friends created and removed restrictions that limited their options was

[1] E-mail: fred.glover@colorado.edu.

PII S0377-2217(97)00259-2

an arresting feature of their approach. This theme was not entirely free of pitfalls, however. In some cases, which always surprised me (since I had the benefit of knowing the "right" answers), some of my friends would adopt a restriction, and then lose track of having done so, failing to realize how their submerged policy was affecting their subsequent search efforts. By neglecting to "undo" the restriction, when it turned out to be wrong, they would never find an optimal answer.

This behavior provided an interesting contrast (and challenge) to Freudian views, which were popular at the time. According to good Freudian theory, self-defeating behavior occurs because of traumatic experiences that cause the suppression of memories which would cause the pain of such experiences to resurface. [2] As a result, affected individuals become incapable of dealing appropriately with related conditions in the present. It was quite clear, however, that the "self-defeating" behavior of my friends had nothing to do with anything remotely traumatic, but was simply a consequence of having so thoroughly (or so long) accepted a particular limitation that the possibility of an alternative was no longer entertained.

Within the Freudian framework, the sources of suppression that block effective behavior are woven into the fabric of social and cultural tabus. The restrictions created by my friends could analogously be viewed as governed by tabus, but evidently of a different character than those envisioned by Freud. The tabus in this case were facilely generated (and then discarded) – and rather than serving as obstacles to thinking, were more usually an integral part of effective problem-solving.

To carry out a problem-solving function, the forms of memory used to manage these tabus departed from the lock-step logic embodied in tree search structures, which AI practitioners were fond of invoking to emulate human problem-solving. [3] The interlinked uses of memories and tabus were evidently susceptible to being handled poorly, yet offered a type of flexibility and vitality that seemed to be worth examining. Consequently, I attempted to capture a simple version of these schemes in the computer program for the assigned AI project. As it turned out, the program performed reasonably well, and in some cases better than my guarded optimism had prepared me to anticipate (though it did not measure up to the performance of a couple of my cleverer friends). This led me to suspect the underlying ideas contained the germ of an approach to go usefully beyond problem-solving designs that were then in vogue in artificial intelligence and optimization. Accordingly, I made a mental note to come back and explore the possibilities more fully one day. [4]

2. Traditional connections

Within the realm of more traditional usage, the word "tabu" has other qualities that make it fitting in reference to tabu search. "Tabu" comes from Tongan, a language of Polynesia, where it was used by the aborigines of Tonga island to indicate things that cannot be touched because they are sacred. According to Webster's Dictionary, the word now also means "a prohibition imposed by social custom as a protective measure" or of something "banned as constituting a risk". These more pragmatic senses of the word accord well with the theme of tabu search. The risk to be avoided in this case is that of following a counter-productive course, including one which may lead to entrapment without hope of escape. On the other hand, as in the broader social context where "protective

[2] This evidently oversimplifies Freud's views. In his book *The Psychopathology of Everyday Life*, Freud acknowledged that not all sources of misguided behavior stemmed from traumatic events, though he still envisioned error and persistent misdirection to derive from painful experiences.

[3] In common with developments by the OR community in branch and bound, the AI community has continued to evolve methods based on variations of tree structures (e.g., such as A* search).

[4] Of course, tabu search has been the outcome of many influences beyond those of this early experiment – including formal ones, provided by relaxation methods and cutting plane procedures in optimization. The area today owes a debt to contributions made by many individuals.

prohibitions" are capable of being superseded when the occasion demands, the "tabus" of tabu search are to be overruled when evidence of a preferred alternative becomes compelling.

The most important association with traditional usage, however, stems from the fact that tabus as normally conceived are transmitted by means of a social memory which is subject to modification over time. This creates the fundamental link to the meaning of "tabu" in tabu search. [5] The forbidden elements of tabu search receive their status by reliance on an evolving memory, which allows this status to shift according to time and circumstance. In common with the behavior evidenced by my friends in the AI problem solving experiment, this memory is highly adaptive in nature. The adaptive memory feature of tabu search undeniably is its most significant aspect.

3. Tabu or not tabu

There are apparent drawbacks, however, in giving a method a name that tends to focus attention on a particular subset of its ideas to the exclusion of others. The temptation to adopt such a focus has sometimes proved too strong to resist. The narrower connotations of "tabu" have prompted a number of researchers to embrace a mind-set that overlooks important characteristics of tabu search. More than a few papers in the literature, for example, examine only a small portion of the elements of short term memory, and examine little or nothing at all of longer term memory. Unfortunately, these papers also often present themselves as embodying the essence of tabu search.

[5] Language takes a curious role here. Just as we may think of using "tabus" to change our relative emphasis on different choices, we may think of altering this emphasis – in exactly the same way – by aspirations. Tabus and aspirations are in a sense mirror images of each other, and we may choose to focus our descriptions in relation to either. Yet, procedurally, the balance is more subtle, and depends on the mechanisms that are convenient for modifying choices in useful ways. For humans and machines alike, there are advantages to considering tabus and aspirations to be more than simple mirror images.

Considering the varied uses of memory that humans employ, the tendency to restrict the treatment of tabu search in this way seems akin to an effort to solve problems by "shutting down" part of the brain. A factor that has reinforced the tendency to examine a limited part of tabu search (aside from convenience, which can be sensible in early stages of an investigation), is that such a focus has sometimes produced very appealing results. When reasonably decent outcomes can be found without great effort, the motive to look further is diminished. The danger, of course, lies in failing to discover significant gains that are likely to be achieved by a more complete approach.

It is appropriate to acknowledge that attention may be given to a limited subset of ideas from an overall search framework for the following reasons:

1. such a focus may help to uncover a better form for the strategies associated with this subset;
2. weaknesses of this subset, when studied in isolation from other ideas, may stand out more clearly – thus yielding insights into the features of a more complete approach that are required to produce a better method;
3. for methods which are susceptible to highly "modular" implementations, as typically occurs for tabu search, simpler designs can readily be made a part of more complex designs.

There remains a consideration that is often overlooked. As is true of metaheuristics in general, tabu search offers a framework for problem solving, as opposed to a rigidly detailed collection of prescriptions about how this framework is best applied. Current research is disclosing instances of this framework that yield remarkably effective outcomes, and is laying the foundation to identify which components of tabu search will ultimately operate to greatest advantage under various conditions. The reference to a "more complete" form of tabu search is therefore partly ambiguous, because the framework includes more options than are appropriate to examine in any single design – at least, given the present stage of our understanding. Nevertheless, there remains a substantial difference between the comprehensiveness of various tabu search procedures that appear in the literature. Notably, the more comprehensive

ones almost always perform appreciably better than the others.

Finally, in the quest to develop advanced tabu search approaches, it is worth noting that there is a difference between "advanced" and "complex". An advanced method may be relatively simple in its structure, and be easy for others to use (as by a self-calibration system for setting its parameters). The difference between rudimentary and advanced methods is analogous to the difference between rudimentary and advanced search paths. A more advanced path may represent a collection of links that are fabricated from a richer set of components. But the path itself need not be complex. In fact, the use of a richer set of components may be the key to allowing the path to have a simple form. An advanced tabu search method may require more effort to identify, yet may be easier to apply than methods which rest on simpler foundations. An important challenge for research is to identify the particular elements that combine to make an advanced method effective – and at the same time, to determine the simplifications that become possible by self-calibrating designs for interrelating these elements. [6]

4. The current special volume on tabu search

In spite of the typical difference in efficacy between different "levels" of tabu search, the papers of this volume cover a range of such levels, from simpler to more encompassing. Partly this stems from the reasons mentioned earlier for justifying a focus on less complete forms of the method. In addition, some of the papers involving simpler variants break new ground by examining areas where tabu search has not yet been applied, and demonstrate the effectiveness of their approaches in obtaining outcomes not matched by methods previously implemented. (These papers provide an invitation to examine such areas in greater depth.) The inclusion of different levels of tabu search implementations also may help the reader to understand the method more fully by disclosing the variation among these levels – as well as the variation among implementations at a given level. [7]

The issue of less complete versus more complete versions of tabu search also arises in studies of "hybrid" methods reported in this volume. In the applications included of this type, improvements are obtained over other metaheuristics by combining them with a tabu search component, and gains are also sometimes achieved over the stand-alone tabu search methods embedded in these hybrids. The tabu search components in these particular studies exclude elements that have proved valuable in other settings, which invite future investigations to determine more thoroughly the potential contributions of alternative strategies. An interesting question is the effect that may be achieved by incorporating such excluded tabu search elements, by comparison to the outcomes derived from substituting mechanisms from other procedures to perform related functions. The topic of hybrid approaches is an exceedingly intriguing one, and merits deeper exploration. The basic issue, however, should not be to simply create a hybrid for its own sake, but to identify the principles that yield the best strategies for fulfilling specific problem solving goals.

The fundamental message is that a great deal remains to be learned about tabu search. Evidently, we also still know very little about how we ourselves use memory in our problem solving. (It is not inconceivable that discoveries about effective uses of memory within our search methods will provide clues about strategies that humans are adept at employing – or may advantageously be

[6] An interesting alternative to developing a self-calibrating design is provided by a software system for optimizing simulation called *OptQuest*, based on tabu search and an associated evolutionary approach called scatter search. For example, this system has been used to fine tune the parameters of a tabu search method for a telecommunications bandwidth packing problem to produce results substantially superior to the original implementation of this procedure reported in the literature. See "OptQuest" on Laguna's Web page: http://spot.colorado.edu/~laguna/Home.html.

[7] A comprehensive treatment of the different levels of tabu search and their implications for solving problems in a variety of settings is given in the book *Tabu Search*, by Glover and Laguna, Kluwer Academic Publishing, 1997.

taught to employ. The potential links between the areas of heuristic search and psychology have scarcely been examined). Unquestionably, in the realm of optimization, we have not yet investigated the strategic possibilities at a level that comes close to disclosing their full potential. The successes of the varied studies reported in this volume provide encouragement that such issues are profitable to probe more fully.