

Difficulty Rating of Sudoku Puzzles by a Computational Model*

Radek Pelánek

Masaryk University Brno, Czech Republic

Abstract

We discuss and evaluate metrics for difficulty rating of Sudoku puzzles. The correlation coefficient with human performance for our best metric is 0.95. The data on human performance were obtained from three web portals and they comprise thousands of hours of human solving over 2000 problems. We provide a simple computational model of human solving activity and evaluate it over collected data. Using the model we show that there are two sources of problem difficulty: complexity of individual steps (logic operations) and structure of dependency among steps. Beside providing a very good Sudoku-tuned metric, we also discuss a metric with few Sudoku-specific details, which still provides good results (correlation coefficient is 0.88). Hence we believe that the approach should be applicable to difficulty rating of other constraint satisfaction problems.

Introduction

The general theme of this work is human problem solving (Simon and Newell 1972). Particularly, we focus on the study of problem difficulty: What determines which problems are difficult for humans? Beside giving us insight into human cognition and thinking, the study of this issue has important applications in human-computer collaboration and training of problem solving skills, e.g., in development of intelligent tutoring systems (Anderson, Boyle, and Reiser 1985; Caine and Cohen 2007).

Difficulty of Problem Solving

We study problem difficulty of one particular problem – Sudoku puzzle. Our specific goal is the following: *Provide a difficulty rating metric for Sudoku puzzle, that achieves as high correlation with human performance (measured by time) as possible.* This goal has direct applications – such metrics are heavily used since Sudoku is currently very popular and even commercially important, and difficulty rating of puzzles is one of the key things which influence user's experience of puzzle solving.

Despite the straightforwardness of our goal and its direct applicability, there is no easily applicable theory that could be used to guide the development of difficulty rating metrics.

*This work is supported by GA ČR grant no. P202/10/0334. Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Currently used Sudoku metrics are usually built in an ad-hoc manner, they are not properly evaluated and their merits are not clear. In general there has been only little research dealing with the issue of problem difficulty; results are available only for few specific puzzles, e.g., Tower of Hanoi and its isomorphs (Kotovsky, Hayes, and Simon 1985), Chinese rings (Kotovsky and Simon 1990), 15-puzzle (Pizlo and Li 2005), and Sokoban puzzle (Jarušek and Pelánek 2010).

The aim of this work goes beyond the specific study of Sudoku puzzle. We would like to raise the interest in the study of problem difficulty, for example by showing that extensive and robust data for study are easily available on the Internet. In this way we would like to contribute towards a theory of difficulty in human problem solving.

Sudoku and Constraint Satisfaction Problems

Sudoku is a well-known number placement puzzle: for a partially filled 9×9 grid, the goal is to place numbers 1 to 9 to each cell in such a way that in each row, column, and 3×3 sub-grid, each number occurs exactly once. Sudoku has been subject to many research studies, particularly with respect to its mathematical and algorithmic properties, e.g., enumerating possible Sudoku grid (Felgenhauer and Jarvis 2005), NP-completeness of generalized version of Sudoku (Yato and Seta 2003), use of constraint propagation (Simonis 2005) or genetic algorithms (Mantere and Koljonen 2007) for solving the puzzle, or algorithms for generating puzzles (O'Sullivan and Horan 2007; Fowler 2009). Recently, also psychological aspects of the puzzle has been studied (Lee, Goodwin, and Johnson-Laird 2008).

We focus on the Sudoku puzzle for several reasons. The Sudoku puzzle has very simple rules, which makes it amenable to analysis. Thanks to its current popularity we can easily obtain large scale data on human solving activity. Sudoku is also a member of an important class of constraint satisfaction problems, which contains many other puzzles and also many real life problems (e.g., timetabling, scheduling). Although we use data for the Sudoku puzzle, our goal is to make the analysis and difficulty metrics as general as possible, so that the results are potentially applicable to other constraint satisfaction problems.

Data on Human Problem Solving

Difficulty rating of Sudoku puzzles is, of course, not a novel problem. The issue of Sudoku difficulty rating is widely discussed among Sudoku players and developers, but it has not been subject to serious scientific evaluation. Current rating algorithms are based mainly on personal experiences and ad-hoc tuning. There are several research papers which discuss methods for difficulty rating (Simonis 2005; Mantere and Koljonen 2007; Henz and Truong 2009); however, these works study the correlation of proposed metric with the difficulty rating provided by the puzzle source (usually a newspaper), not with the data on human performance. Such analysis is not very meaningful since the rating provided in puzzle sources is just another rating provided by a computer program (nearly all published puzzles are generated and rated by a computer). The only work that we are aware of and that uses data on real human performance is the brief report by Leone, Mills, and Vaswani (2008).

Due to the popularity of the Sudoku puzzle we have been able to obtain data capturing hundreds of thousands of hours of human problem solving activity (approximately 2000 puzzles, hundreds of human solvers for each puzzle). This means that we have data several orders of magnitude more extensive than the usual data used in study of human problem solving – most previous research is based on tens or hundreds of hours of human problem solving activity (usually about 20 people and 5 puzzles). Even though this way of data collection has its disadvantages (e.g., lack of direct control over participants), thanks to the scale of the “experiment”, the data are robust and applicable for research purposes.

Sudoku and Constraint Satisfaction Problems

The Sudoku puzzle is a special case of a more general type of problems called constraint satisfaction problems (CSP). In this section we describe both the general CSP and the specific Sudoku problem. We also discuss basic techniques for solving these problems.

Constraint Satisfaction Problems

Constraint satisfaction problem is given by a set of variables $X = \{x_1, \dots, x_n\}$, a set of domains of variables $\{D_1, \dots, D_n\}$ (we consider only finite domains), and a set of constraints $\{C_1, \dots, C_m\}$. Each constraint involves some subset of variables and specifies allowed combinations of variable values (usually given in a symbolic form, e.g., $x_1 \neq x_2$).

A solution of a constraint satisfaction problem is an assignment of values to all variables such that all constraints are satisfied. The class of CSPs contains many puzzles (e.g., eight queen problem, cryptarithmic puzzle) as well as many important practical problems (map coloring problems, timetabling problems, transportation scheduling). The general CSP is NP-complete.

Sudoku Puzzle

Sudoku puzzle is a grid of 9×9 cells, which are divided into nine 3×3 sub-grids, partially filled with numbers 1 to 9.

1	2,4	2,4	2,3,4
2	2,4	1	2,3,4
4	1,2	3	1,2
2	3	2,4	1,2,4

Figure 1: A sample 4×4 Sudoku puzzle with enumerated candidate sets. Circle marks naked single, rectangle marks hidden single.

The solution of the puzzle is a complete assignment of numbers 1 to 9 to all cells in the grid such that each row, column and sub-grid contains exactly one occurrence of each number from the set $\{1, \dots, 9\}$. Sudoku puzzle is well posed, if it admits exactly one solution. We study only well-posed puzzles.

Sudoku puzzle can be easily generalized for any grid size of $n^2 \times n^2$ and values from 1 to n^2 (Simonis 2005). Moreover, there are many variants of Sudoku which use non-regular sub-grids (e.g. pentomino), or additional constraint (e.g., arithmetic comparison of values). In this work we use 4×4 Sudoku puzzle as a small running example, otherwise we consider solely the classical 9×9 Sudoku puzzles. Sudoku can be easily expressed as a constraint satisfaction problem (Simonis 2005).

Constraint Propagation

Classical Sudoku can be easily solved computationally by backtracking search (Pelánek 2011). However, backtracking search is not very relevant to estimating human performance. Another approach to solving CSP is constraint propagation. This method tries to find values of (some) variables by reasoning about constraints. For each variable x_i we define a current candidate set – a set of such values that do not lead to direct violation of any constraint (see Figure 1). By reasoning about candidate sets and constraints we can often derive solution without any search.

Constraint propagation is not guaranteed to find a solution, but it may be more efficient than backtracking search and can also be combined with backtracking search to produce superior results. We are interested in constraint propagation particularly because this is the natural way how humans try to solve CSPs.

Let us consider specifically the case of Sudoku puzzle. Human solving of Sudoku proceeds by sequence of steps in which values are filled into cells. Two basic techniques directly correspond to the rules of the puzzle (see also Figure 1). *Naked single technique* (also called singleton, single value, forced value, exclusion principle): For a given cell there is only one value that can go into the cell, because all other values occur in row, column or sub-grid of the cell (any other number would lead to a direct violation of rules). *Hidden single technique* (also called naked value, inclusion principle): For a given unit (row, column or sub-grid) there exists only one cell which can contain a given value (all other placements would lead to a direct violation of rules).

Sudoku problems solvable by iteration of these two techniques are further denoted as “simple Sudoku”. Most of the publicly used puzzles which are ranked as easy or mild are simple Sudokus. There exists many advanced techniques, such as pointing, chaining, naked and hidden pairs (see, e.g., (Juillerat 2009)), but we do not elaborate on these techniques in order to keep Sudoku-specific details minimized.

Data on Human Sudoku Solving

For obtaining data on human problem solving we exploited the current popularity of on-line puzzle solving and particularly the popularity of Sudoku puzzle. The data were obtained from three Sudoku web portals. From the server `sudoku.org.uk` we have summary data provided by the server: total number of solvers (the mean is 1307 solvers per puzzle) and the mean time to solve the puzzle (no data on individual solvers), we have used 731 puzzles. From the portal `fed-sudoku.eu` we have also information about individual users (the overall time to solve each puzzle). We have used 1088 puzzles, the mean number of solvers is 131 per puzzle. Finally, from the server `czech-sudoku.com` we have not just the time to solve the puzzle, but also the record of each play. We analyzed these detailed records for 60 users and 15 puzzles.

As a measure of problem difficulty for humans we use the mean solution time from all solutions. We have thoroughly checked that the main results are not dependent on this choice (we have done the analysis also for median time or for mean time computed only from a selected subset of active users). For detailed analysis and description of the datasets see (Pelánek 2011).

Computational Model of Human Solver

In this section we discuss a simple model of a human Sudoku solving. We also provide evaluation of the model using the data on human problem solving. Although we specify and evaluate the model only for Sudoku puzzle, the basic model is general, easily modifiable and applicable to other CSPs.

Our main motivation for developing the model is difficulty rating. Nevertheless, the model could be useful in other applications as well, e.g., as a part of a tutoring system (Caine and Cohen 2007) or for detection of cheating in Internet Sudoku competitions (if the user fills repeatedly cells in wrong order, then it is probable that he did use computer solver to solve the puzzle).

General Model

We propose a simple model of human CSP solving, which is based on the following assumptions¹. Humans are not good at performing systematic search, and there are not willing to do so. Humans rather try to solve CSPs by ‘logic techniques’, i.e., by constraint propagation. Moreover humans prefer ‘simple’ techniques over ‘difficult’ ones (we elaborate on difficulty of logic techniques below).

¹We are not aware of any scientific research which could be used to support these assumptions, but there is ample support for them in popular books about puzzle solving.

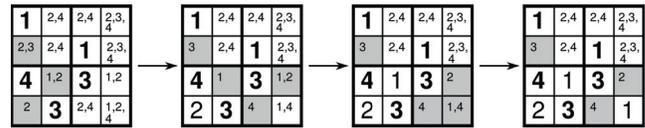


Figure 2: Example of a model run on a sample 4×4 Sudoku puzzle. Grey cells are cells for which the value can be directly determined using one of the simple techniques (naked single, hidden single). In each step one of these cells is selected randomly. Only first three steps of the model run are shown.

The model proceeds by repeatedly executing the following steps until the problem is solved (see Figure 2 for illustration):

1. Let L be the simplest logic technique which yields for the current state some result (variable assignment, restriction of a candidate set).
2. Let a by an action which can be performed by the technique L . If there are several possibilities how to apply L in the current state, select one of them randomly.
3. Apply a and obtain new current state.

Note that this model makes two simplifying assumptions: at first that the solver does not make any mistakes (i.e., no need to backtrack) and that the solver is always able to make progress using some logic technique, i.e., the solver does not need to perform search. These assumptions are reasonable for Sudoku puzzle and are supported by our data on human problem solving. For other CSPs it may be necessary to extend the model.

Logic Techniques and Their Difficulty Rating

To specify the stated abstract model, we have to provide list of logic techniques and their difficulty rating. The usual approach used by Sudoku tools is based on a list of logic techniques which are supposed to be simulations of techniques used by humans; each of these techniques is assigned difficulty rating. This rating is provided by the tool developer, usually based on personal experience and common knowledge.

This approach has the disadvantage that it contains lot of ad-hoc parameters and it is highly Sudoku-specific, i.e., it gives us limited insight into human problem solving and it is not portable to other problems (the success of the approach is based on significant experience with the problem).

We propose an alternative approach to classification of logic techniques. The approach is based on the assumption that many advanced logic techniques are in fact “short-cuts” for a search (what-if reasoning).

We therefore provide rating of difficulty of logic techniques with the use of search. This approach contains nearly no parameters and is not specific to Sudoku (i.e., it is applicable to any CSP). The only Sudoku specific issue is the selection and realization of “simple” techniques – in our case these are hidden single and naked single techniques; note that these techniques are basically derived from the rules of

the problem. For most CSP problems it should be possible to derive basic simple techniques on a similar basis.

Let us suppose that we have a state in which the specified simple techniques do not yield any progress. For each unassigned variable (empty cell) we compute a “refutation score”, this score expresses the difficulty of assigning the correct value to this variable in the given state by refuting all other possible candidates.

For each wrong candidate value v we denote ref_v the smallest number of simple steps which are necessary to demonstrate the inconsistency of the assignment. The “ideal refutation score” is obtained as a sum of values ref_v . If some of the values is not refutable by simple steps, we set the score to ∞ .

The computation of ref_v can be done by breadth-first search over possible puzzle states, but it is computationally expensive and anyway the systematic search does not correspond to human behavior. Therefore we use randomized approach analogical to our main model – instead of computing the smallest number of steps necessary to refute a given value, we just use a randomized sequence of simple steps and count the number of steps needed to reach an inconsistency. The refutation score is thus a randomized variable.

The variable (cell) with the lowest score is deemed to be the easiest to fill and the refutation score is used as a difficulty rating of an (unknown) logic technique. For all our considered Sudoku puzzles there was always at least one cell with finite score; for more complex problems it may be necessary to further specify the model for the case that all refutation scores have value ∞ .

Evaluation of the Model

Using the described notions we specify a “Simple Sudoku Solver” (SiSuS) model: the general model described in Section with two hard-wired logic techniques (hidden single, naked single) of equal difficulty which uses refutation score when the basic techniques are not applicable.

We have evaluated the SiSuS model over detailed data records from `czech-sudoku.com`. To evaluate our model we compare the order in which the cells are filled by humans and the model. For the evaluation we used 15 selected puzzles of wide range of difficulty (from very easy to very difficult). Each puzzle was solved by 10 to 60 solvers.

Based on the data records of human solvers we computed the mean order for each cell. Similarly we computed for each cell mean order over 30 randomized runs of our model. In most cases the Spearman’s correlation coefficient between the two orderings is between 0.85 and 0.95; see (Pelánek 2011) for illustrations over specific examples. Best results are obtained for puzzles of intermediate difficulty. For very easy puzzles there are many ways in which cells can be filled and therefore it is hard to predict the exact order (in this cases the order also differs among individual solvers). Difficult puzzles cannot be solved by the basic techniques used by the model and hence the prediction is again bit worse. Nevertheless, given the simplicity of the SiSuS model, we consider the overall performance to be very good.

Difficulty Metrics

Based on the model of human solution progress we now provide several difficulty metrics and evaluate them on the data on human behaviour. For all studied metrics we report the Pearson’s correlation coefficient.

Note that difficulty rating is interwoven with modeling human solvers. Difficulty metrics are based on the data collected by simulating the model of human solver, but the model depends on the rating of difficulty of techniques. Models which incorporate many logic technique depend on the intuition of the human designer (alternatively they could use some kind of bootstrapping).

Combining Rating of Logic Techniques

Given a model of a human solver, a straightforward approach to difficulty rating is to run the model, count how often each logic technique is used and produce the overall rating as a simple function of these statistics. This is the approach used by most Sudoku generators. For our evaluation, we use the following metrics:

Serate metric Default metric used by the Sudoku Explainer tool (Juillerat 2009), which uses more than 20 techniques. The metric is a maximal difficulty of a used logic technique.

Serate LM metric Linear model over techniques used by the Sudoku Explainer tool; this approach is inspired by (Leone, Mills, and Vaswani 2008). We compute how many times each logic technique² was used over each problem. Using half of the problems as a training set we compute parameters for a linear model; the metric is evaluated on the remaining problems (a test set).

Fowler’s metric Default metric used by G. Fowler’s tool (Fowler 2009); the metric is given by a (rather complicated) expression over number of occurrences of each logic techniques (with ad-hoc parameter values).

Refutation sum metric Mean sum of refutation scores over 30 randomized run of our SiSuS model.

Dependency Metric

So far we have focused on the difficulty involved in single steps. The overall organization of these steps was considered only as a simple function of difficulty ratings of individual steps. Insufficiency of this approach can be seen particularly for simple Sudokus – these problems are solvable by the basic simple techniques (i.e., the above describe metrics return very similar numbers), but for humans there are still significant differences in difficulty (some problems are more than two times more difficult than others).

Some of this additional difficulty can be explained by the concept of ‘dependency’ among steps in the solution process (applications of logic techniques). An important aspect of human CSP solving is “the number of possibilities leading to a next step” in each step. For example in our small

²We take into account only techniques which were used in at least 0.5% of all technique applications. There are 13 such techniques, all other techniques were grouped together.

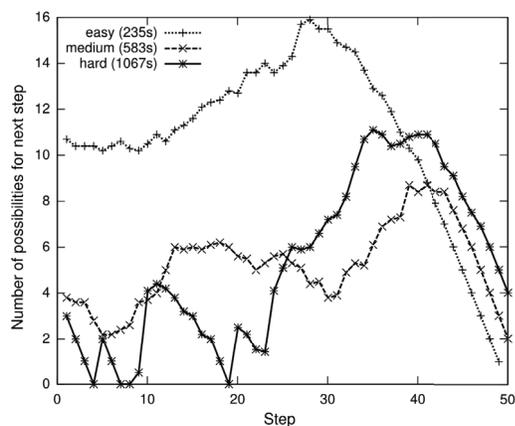


Figure 3: Dependency among steps captured by graph of number of possibilities for the next step. Results for three sample puzzles of different difficulty are shown (the difficulty is indicated by mean solution time of human solvers).

Sudoku example from Figure 2, there are 3 possibilities in the first step, 4 possibilities in the second and third steps, and so on. It is quite clear that for the classical 9×9 Sudoku it makes a big difference if we can in the first step apply a logic technique at 10 different cells or only at just 2.

To apply this idea, we count in each step of the SiSuS model the number of possibilities to apply a simple technique. Since the model is randomized, we run the model several times and compute for each step mean number of possibilities. Figure 3 shows a difference among several specific instances – it shows that for easy problem there are many possibilities for progress in each step whereas for hard problem there are only few of them.

To specify a difficulty metric, we need to convert the graphs in Figure 3 to a single number. We simply compute the mean over the first k steps (k is a parameter of the metric). But what is a good value of k ? As illustrated by examples in the Figure 3, in the second half of the solution there are usually many possibilities for all problems; i.e., these steps probably do not contribute to the difficulty and therefore it is better to limit the parameter k , on the other hand too small k ignores potentially useful information. We have evaluated the preciseness of the metric with respect to the parameter k over our data sets. The results show that a suitable value of k is slightly dependent on the data set, but generally it is between 20 and 30 and results are not too much dependent on the precise choice of k (for the interval 20 to 30).

Evaluation

Except for the metrics described above, we also evaluated combinations of metrics, more specifically linear models over several metrics. Parameters of linear models were determined over a training set (one half of the problems), results were evaluated over the other half of models (testing set). We evaluated two linear models. The first combined metric is based on data obtained only from our SiSuS model

Table 1: Correlation coefficients between metrics and human results. Refutation sum metric is not applicable to simple problems.

metric	fed-sudoku		sudoku.org	
	all	simple	all	simple
number of givens	0.25	0.22	0.27	0.34
Serate	0.70	0.55	0.86	0.28
Serate LM	0.78	0.60	0.86	0.66
Fowler’s	0.68	0.53	0.87	0.64
Refutation sum	0.68	–	0.83	–
Dependency	0.67	0.73	0.69	0.78
Combined (RD)	0.74	–	0.88	–
Combined (SFRD)	0.84	0.75	0.95	0.83

(linear combination of Refutation sum and Dependency metric; denoted “RD” in Table 1). The second combined metric is a based on four metrics (Serate, Fowler’s, Refutation sum, Dependency; denoted “SFRD” in Table 1).

Results are given in Table 1. Figure 4 gives scatter plots for combined metric SFRD as an illustration of the distribution of the data points.

We get consistently better results for `sudoku.org.uk` than for `fed-sudoku.eu`. This is probably mainly due to the wider variability of difficulty in the `sudoku.org.uk` dataset. Beside the difference in absolute numbers, all other below discussed trends are the same over both datasets.

For the “Simple” subset of puzzles (solvable only by hidden single and naked single techniques), previously studied metrics (Serate, Fowler’s) achieve rather poor results; on the other hand, the new Dependency metric works quite well.

The Refutation sum metric achieves only slightly worse results than classical metrics (Serate, Fowler’s), despite the fact that it is much more general and simpler technique with only little Sudoku specific aspects (particularly it does not have ad-hoc parameters).

Serate LM metric (linear model over data about the usage of 14 logic techniques) achieves similar results as basic Serate metric. Fowler’s metric, which differs in details and parameter values but uses the same basic approach as Serate, also achieves similar results. It seems that given the basic approach, the selection of exact parameter values is not that much important. Nevertheless, by combining 4 different metrics, we can significantly improve the overall performance and achieve really good performance of the metric.

Conclusions

Current popularity of puzzle solving via the Internet enables us to easily collect extensive data on human problem solving. Although such data collection is not done under controlled laboratory conditions, our analysis shows that data from the Internet may be robust and definitively useful. In our evaluation we used two different datasets; although we did get different absolute results for each dataset, relative results (comparison among different techniques) was nearly the same – this supports our believe in robustness and usefulness of the data collected from the Internet. In this work we use the data to study and evaluate difficulty ratings of

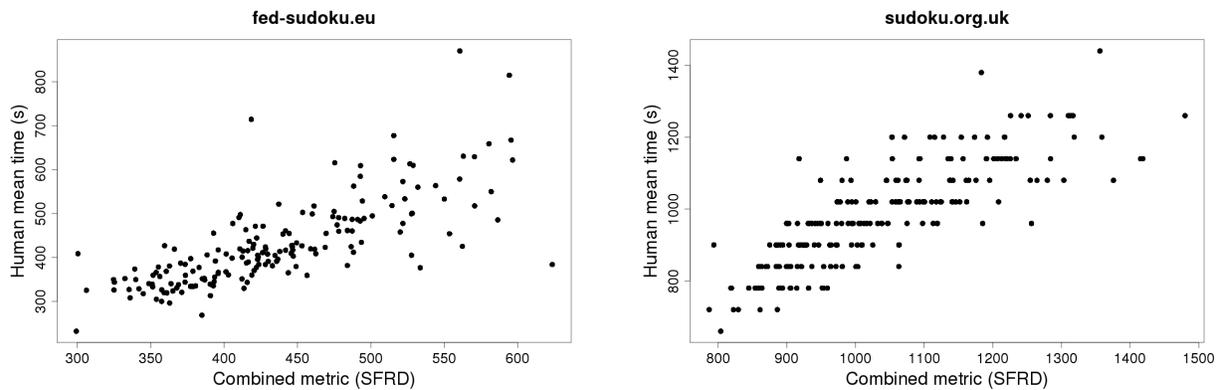


Figure 4: Scatter plots showing relation between prediction of difficulty by combined rating metric and real difficulty; each dot corresponds to one Sudoku puzzle. Graphs correspond to the last line in Table 1.

a sample problem; but the approach could be used also for other problems and for studies of other cognitive issues (e.g., what kind of errors humans do).

In this work we study a Sudoku puzzle as an example of a constraint satisfaction problem. We provide a general model of human CSP solving. We show that by instantiating the model with only few and simple Sudoku-specific details, we can obtain quite reasonable difficulty rating metric (correlation coefficient up to 0.88). By combining several techniques which are specifically tuned for Sudoku we are able to obtain very good difficulty rating metric (correlation coefficient up to 0.95).

We identify two aspect which influence the problem difficulty: difficulty of individual logic steps during the solution and dependency among individual steps. Previously used techniques (Fowler 2009; Juillerat 2009) focused only on individual logic steps. The novel concept of dependency enabled us to significantly improve the performance of rating.

Acknowledgement

The author thanks to webmasters of `fed-sudoku.eu` and `czech-sudoku.com` portals for providing the data and Jiří Šimša and Petr Jarušek for inspiring discussions.

References

- Anderson, J.; Boyle, C.; and Reiser, B. 1985. Intelligent tutoring systems. *Science* 228(4698):456–462.
- Caine, A., and Cohen, R. 2007. Tutoring an entire game with dynamic strategy graphs: The mixed-initiative sudoku tutor. *Journal of Computers* 2(1):20–32.
- Felgenhauer, B., and Jarvis, F. 2005. Enumerating possible Sudoku grids.
- Fowler, G. 2009. A 9x9 sudoku solver and generator. AT&T Labs Research.
- Henz, M., and Truong, H. 2009. Sudoku Sat – A Tool for Analyzing Difficult Sudoku Puzzles. *Tools and Applications with Artificial Intelligence* 25–35.
- Jarušek, P., and Pelánek, R. 2010. Difficulty rating of sokoban puzzle. In *Proc. of the Fifth Starting AI Researchers’ Symposium (STAIRS 2010)*. IOS Press.
- Juillerat, N. 2009. Sudoku Explainer.
- Kotovsky, K., and Simon, H. 1990. What Makes Some Problems Really Hard: Explorations in the Problem Space of Difficulty. *Cognitive Psychology* 22(2):143–83.
- Kotovsky, K.; Hayes, J.; and Simon, H. 1985. Why are some problems hard? Evidence from tower of Hanoi. *Cognitive psychology* 17(2):248–294.
- Lee, N.; Goodwin, G.; and Johnson-Laird, P. 2008. The psychological puzzle of Sudoku. *Thinking & Reasoning* 14(4):342–364.
- Leone, A.; Mills, D.; and Vaswani, P. 2008. Sudoku: Bagging a difficulty metric and building up puzzles. *Mathematical Contest in Modeling*, University of Washington.
- Mantere, T., and Koljonen, J. 2007. Solving, rating and generating Sudoku puzzles with GA. In *IEEE Congress on Evolutionary Computation (CEC 2007)*, 1382–1389.
- O’Sullivan, B., and Horan, J. 2007. Generating and Solving Logic Puzzles through Constraint Satisfaction. In *Proc. of the 22nd national conference on Artificial intelligence*, volume 2, 1974–1975. AAAI Press.
- Pelánek, R. 2011. Human problem solving: Sudoku case study. Technical Report FIMU-RS-2011-01, Masaryk University Brno.
- Pizlo, Z., and Li, Z. 2005. Solving combinatorial problems: The 15-puzzle. *Memory and Cognition* 33(6):1069.
- Simon, H., and Newell, A. 1972. *Human problem solving*. Prentice Hall.
- Simonis, H. 2005. Sudoku as a constraint problem. In *Proc. 4th Int. Workshop on Modelling and Reformulating Constraint Satisfaction Problems*, 13–27.
- Yato, T., and Seta, T. 2003. Complexity and completeness of finding another solution and its application to puzzles. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* 86(5):1052–1060.