

Measures of Intrinsic Hardness for Constraint Satisfaction Problem Instances

George Boukeas, Constantinos Halatsis,
Vassilis Zissimopoulos, and Panagiotis Stamatopoulos

Department of Informatics and Telecommunications, University of Athens

Abstract. Our aim is to investigate the factors which determine the intrinsic hardness of constructing a solution to any particular constraint satisfaction problem instance, regardless of the algorithm employed. The line of reasoning is roughly the following: There exists a set of distinct, possibly overlapping, trajectories through the states of the search space, which start at the unique initial state and terminate at complete feasible assignments. These trajectories are named solution paths. The entropy of the distribution of solution paths among the states of each level of the search space provides a measure of the amount of choice available for selecting a solution path at that level. This measure of choice is named solution path diversity. Intrinsic instance hardness is identified with the deficit in solution path diversity and is shown to be linked to the distribution of instance solutions as well as constrainedness, an established hardness measure.

1 Introduction

A constraint satisfaction problem consists of a set of variables and a set of constraints. A variable which has been given a value is said to be *instantiated*. A set of instantiations to i distinct variables is an *assignment* α_i of size i . If the size of an assignment α_n equals the number of problem variables n then it is a *complete assignment*. A complete assignment which satisfies all problem constraints is a feasible *solution*. The set of solutions to a problem instance is denoted by \mathcal{S} . Given a constraint satisfaction problem instance, the goal is to find a feasible solution or to prove that none exists. In order to accomplish this, *constructive search methods* start from the empty assignment and iteratively extend partial assignments until a feasible solution is found. Therefore, the complete search space comprises all distinct assignments of all sizes, partitioned into disjoint levels according to size. An extensive presentation can be found in [1]. In contrast, *repair search methods* iteratively transform complete assignments until a feasible solution is found. Therefore, the complete search space comprises only complete assignments.

This research aims at investigating the factors which determine the *intrinsic hardness* of constructing a solution to any particular constraint satisfaction problem instance. Other than the assumption that a constructive algorithm is employed, instance hardness is treated in a manner independent of the particular

tree-search method used, hence the use of the term “intrinsic”. In line with [1], as well as [2], [3], our viewpoint focuses on the structure of the induced search space, which allows an abstraction away from problem-specific properties.

The line of reasoning is roughly the following: There exists a set of distinct, possibly overlapping, trajectories through the states of the search space, which start at the unique initial state (the empty assignment containing no instantiations) and terminate at complete feasible assignments. These trajectories are named *solution paths*. See Fig. 1 for an illustration of the search space and solution paths. The entropy of a distribution is a measure of *choice* in selecting




Fig. 1. The search space of a constructive method for instances with $n = 3$ binary variables. The search space is partitioned into disjoint levels, with each level i containing the $2^i C(n, i)$ possible instantiations of size i . The $n!$ possible paths to a particular complete assignment are also depicted. The search space for a repair method comprises the complete assignments in the shaded rectangle

an event [4]. The entropy of the distribution of solution paths among the states of each level of the search space provides a measure of the amount of choice available for selecting a solution path at that level. This measure of choice is named *solution path diversity*. Intrinsic instance hardness is identified with the *deficit* in solution path diversity: the lower the amount of choice available to any algorithm for selecting a solution path, the higher the intrinsic instance hardness. Choice is inherent in the search space of the problem instance and thus independent of the algorithm used to traverse it. In this work, the focus lies on investigating the choices offered by the search space of instances, not on how particular algorithms may make use of such choices. Therefore, throughout the presentation, the reader should bear in mind that the notion of intrinsic instance hardness is not to be confused with computational cost. The former is invariant for a problem instance and manifests itself through the latter, the exact value of which depends on the particular algorithm employed.

To our knowledge, the application of such reasoning in order to characterize intrinsic instance hardness is novel. Moreover, it uncovers a remarkable link between intrinsic instance hardness and *constrainedness* [2]. The latter has been successfully introduced as an estimator of solution cost for problem ensembles

but its justification as *average information content per variable* in [5] is intuitive rather than theoretical, especially since it provides no link between constrainedness and computational cost. Here, it is shown that the constrainedness of an instance can be regarded as *deficit in the total amount of choice available for selecting a solution path, averaged over all levels of the search tree*.

There is substantial relevance between the object of this work and research in phase transitions as well as other areas. Similarities and deviations are elaborated upon in Sect. 2. Apart from that, measures of choice are discussed in Sect. 3, leading to measures of intrinsic instance hardness in Sect. 4. The presentation is concluded in Sect. 5 which underlines the contributions of this work. The majority of the propositions to be found herein have short proofs, mostly involving algebraic manipulations, and have been omitted due to space restrictions.

2 Related Research

Research in *phase transitions* investigates how problem structure affects the average computational cost [6]. Although such work deals with *ensembles of instances* rather than individual instances, our research is intimately linked to [2] and [1]. In the former, the *number* of feasible solutions is used in the definition of *constrainedness*, an established predictor of computational cost. In the latter, the *number* of partial consistent assignments (goods) at each level of the search space is also employed. Refining these approaches, we use information-theoretic measures, similar to those found in [7], to quantify the *distribution* of solutions and, consequently, partial consistent assignments. This uncovers an interesting link between the measures proposed here and constrainedness. There is also a close connection between our work and [8], which investigates the relationship between algorithmic complexity and computational cost.

3 Measures of Choice

3.1 Solution Path Diversity

There are $n!$ paths towards any solution $\alpha_n \in \mathcal{S}$, corresponding to the $n!$ distinct orderings of the instantiations in α_n . The total number of paths towards the solutions in \mathcal{S} is thus $n!|\mathcal{S}|$. The subset of these solution paths containing a particular assignment α_i is $i!(n-i)!\rho_{\mathcal{S}}(\alpha_i)$, where $\rho_{\mathcal{S}}(\alpha_i)$ is the number of complete assignments which are solutions of \mathcal{S} and are reachable from α_i . Consequently, the fraction $\wp(\alpha_i)$ of solution paths containing partial assignment α_i is:

$$\wp(\alpha_i) = \frac{\rho_{\mathcal{S}}(\alpha_i)}{|\mathcal{S}|C(n, i)}$$

where $C(n, i) = n!/(i!(n-i)!)$. These fractions are not probabilities, although they can be interpreted as such. Essentially, the fractions $\wp(\alpha_i)$ define the *distribution* of solution paths among the assignments of level i . Apparently, the *entropy* of this distribution provides a measure of choice in selecting a solution path at level i .

Definition 1 (\mathcal{D}_i). *The entropy of the distribution of solution paths among the assignments of level i , as defined by the fractions $\wp(\alpha_i)$, is called solution path diversity at level i :*

$$\mathcal{D}_i = - \sum_{\alpha_i} \wp(\alpha_i) \log \wp(\alpha_i)$$

Solution path diversity is a concise quantitative measure of the amount of choice available to any constructive algorithm for selecting a solution path at level i . It is measured in *bits* (assuming the logarithm is base 2). A more refined measure of choice can be obtained through the fraction $\wp(\alpha_{i-1} \cdot \alpha_i)$ of solution paths containing consecutive assignments α_{i-1} and α_i , which can be shown to be:

$$\wp(\alpha_{i-1} \cdot \alpha_i) = \frac{\sigma(\alpha_{i-1}, \alpha_i) \rho_S(\alpha_i)}{i|\mathcal{S}| C(n, i)}$$

The fractions $\wp(\alpha_{i-1} \cdot \alpha_i)$ define the distribution of solution paths among the instantiations leading to level i .

Definition 2 ($\Delta\mathcal{D}_i$). *The conditional entropy of the distribution of solution paths among the instantiations leading to level i , as defined by the fractions $\wp(\alpha_{i-1} \cdot \alpha_i)$, is called conditional solution path diversity at level i :*

$$\Delta\mathcal{D}_i = - \sum_{\alpha_{i-1}} \sum_{\alpha_i} \wp(\alpha_{i-1} \cdot \alpha_i) \log \frac{\wp(\alpha_{i-1} \cdot \alpha_i)}{\wp(\alpha_{i-1})}$$

Conditional solution path diversity is a measure of the amount of choice available to any constructive algorithm for selecting a solution path at level i , having reached level $i - 1$. It is a monotonic decreasing function with respect to i and is measured in *bits per level*. Figure 2 depicts the search spaces for two constraint satisfaction problem instances with three binary variables and two solutions, along with the distribution of solution paths. It is straightforward to show that the two forms of solution path diversity are connected in the following manner:

$$\Delta\mathcal{D}_i = \mathcal{D}_i - \mathcal{D}_{i-1} + \log i \tag{1}$$

Recall that the search space involves all possible partial assignments and therefore (implicitly) all possible variable orderings for constructing these assignments. It can be shown that this is the reason for the appearance of the $\log i$ term.

3.2 The Distribution of Solutions

Assume the n problem variables are all binary. There exist 2^n distinct complete assignments which correspond to the 2^n vertices of a n -dimensional hypercube. Let the vertices which correspond to solutions be colored black and let the rest of the vertices be colored white. This coloring, which essentially reflects the placement of solutions on the vertices of the hypercube, is the *distribution of*




Fig. 2. The search space and solution paths for instances with $n = 3$ binary variables and $|\mathcal{S}| = 2$ solutions, with the distance between the solutions being (a) minimal: one instantiation and (b) maximal: n instantiations. Circle area and line thickness reflect the distribution of solution paths among assignments and instantiations




Fig. 3. The distribution of solutions for two non-isomorphic instances of three binary variables and two solutions. The instances correspond to the search spaces of Fig. 2

solutions. As an example, Fig. 3 contains the distribution of solutions for two problem instances with three binary variables and two solutions.

Problem instances are *isomorphic* if the corresponding distributions of solutions (colored hypercubes) are isomorphic. If the solution sets \mathcal{S} and \mathcal{T} of two distinct instances of n variables are isomorphic, then for every level i of the search space, it holds that:

$$\mathcal{D}_i^{\mathcal{S}} = \mathcal{D}_i^{\mathcal{T}}$$

There is currently no available proof for the inverse (that equal entropy implies isomorphism) although it is our strong belief that it is also true. This entails that solution path diversity \mathcal{D}_i directly reflects the distribution of solutions and is both necessary and sufficient for discriminating between instances with non-isomorphic solution sets. It is not uncommon for entropy to be an invariant of isomorphic structures (see [9]).

3.3 The Sum Property

An accumulated measure of choice is obtained by summing (1) over all levels i of the search space. This is acceptable since the conditional solution path diversity $\Delta\mathcal{D}_i$ pertains exclusively to level i of the search space.

Proposition 1 (\mathbf{D}^n). *The sum of the conditional solution path diversities $\Delta\mathcal{D}_i$ over all levels i of the search space is the total solution path diversity \mathbf{D}^n and is identical among all problem instances with the same number of variables n and the same number of solutions $|\mathcal{S}|$.*

$$\mathbf{D}^n(\mathcal{S}) = \sum_{i=1}^n \Delta\mathcal{D}_i = \log(n!|\mathcal{S}|)$$

Not surprisingly, $\mathbf{D}^n(\mathcal{S})$ depends upon $n!|\mathcal{S}|$, the number of solution paths. Proposition 1 holds even though the individual $\Delta\mathcal{D}_i$'s which make up \mathbf{D}^n need not be identical among all such instances. This will be referred to as *the sum property*.

3.4 General Path Diversity

Solution path diversity cannot, by itself, serve as a measure of hardness because it is a *relevant* quantity. It measures the amount of choice available in selecting a solution path and therefore becomes meaningful only when compared to the total amount of choice available in selecting *any* path. This is computed using the entropy of the uniform distribution of all possible paths among all possible assignments. For the sake of simplicity, it will be assumed that variables are binary.

Definition 3 (\mathcal{G}_i). *The entropy of the uniform distribution of the $n!2^n$ distinct possible paths among the $2^i C(n, i)$ distinct possible assignments of level i is called general path diversity at level i :*

$$\mathcal{G}_i = \log [2^i C(n, i)] = i + \log C(n, i)$$

General path diversity is a measure of the amount of choice available in selecting any path at level i . It essentially reflects the size of the search space.

Definition 4 ($\Delta\mathcal{G}_i$). *The conditional entropy of the uniform distribution of the $n!2^n$ distinct possible paths among the $C(n, i)2^{2i}$ distinct possible instantiations leading to level i is called conditional general path diversity at level i .*

$$\Delta\mathcal{G}_i = \mathcal{G}_i - \mathcal{G}_{i-1} + \log i = 1 + \log(n - i + 1)$$

Conditional general path diversity is a measure of the amount of choice available in selecting any path at level i , having reached level $i - 1$. It is a monotonic decreasing function with respect to i .

In a sense, the (conditional) general path diversity defines the *maximal attainable value* for the (conditional) solution path diversity. If all paths of a search space were solution paths, then the solution path diversity would be equal to the general path diversity.

Proposition 2 (\mathbf{G}^n). *The sum of the conditional general path diversities $\Delta\mathcal{G}_i$ over all levels i of the search space is the total general path diversity \mathbf{G}^n and is identical among all problem instances with the same number of variables n .*

$$\mathbf{G}^n = \sum_{i=1}^n \Delta\mathcal{G}_i = \log(n!2^n)$$

4 Measures of Intrinsic Hardness

The easiest possible problem instance arises when all assignments are feasible and every path is thus a solution path. In this case, the solution path diversity \mathcal{D}_i is maximized and equal to the general path diversity \mathcal{G}_i at every level i . In every other case, there is a *deficit* in the amount of choice available for selecting a solution path. The magnitude of this deficit identifies intrinsic instance hardness.

Definition 5 (\mathcal{H}_i). *The difference between the general path diversity \mathcal{G}_i and the solution path diversity \mathcal{D}_i is called the intrinsic instance hardness at level i .*

$$\mathcal{H}_i = \mathcal{G}_i - \mathcal{D}_i$$

Definition 6 ($\Delta\mathcal{H}_i$). *The difference between the conditional general path diversity $\Delta\mathcal{G}_i$ and the conditional solution path diversity $\Delta\mathcal{D}_i$ is called the conditional intrinsic instance hardness at level i :*

$$\Delta\mathcal{H}_i = \Delta\mathcal{G}_i - \Delta\mathcal{D}_i = \mathcal{H}_i - \mathcal{H}_{i-1}$$

The intrinsic instance hardness \mathcal{H}_i and $\Delta\mathcal{H}_i$ rises as the solution path diversity \mathcal{D}_i and $\Delta\mathcal{D}_i$ drops. It is maximal when there are no solutions and minimal when there are no infeasible assignments. It also rises along with the size of the search space, which manifests itself through the general path diversity \mathcal{G}_i and $\Delta\mathcal{G}_i$. An illustration can be found in Fig. 4.

Both forms \mathcal{H}_i and $\Delta\mathcal{H}_i$ of intrinsic instance hardness are monotonic increasing functions with respect to the search level i , which is due to the fact that solution paths tend to spread at deeper levels of the search space. This is not in conflict with the intuition that subproblems encountered by search algorithms become easier as the search progresses into deeper levels. Both \mathcal{H}_i and $\Delta\mathcal{H}_i$ involve *averages* over all the assignments or instantiations of a level and thus contain no information about the hardness of the subproblems that may be encountered *during search* by a particular search algorithm (a subject of vast practical significance discussed in [10], [5]).

4.1 The Sum Property and Constrainedness

Using Props. 1 and 2, it follows that the sum property also holds for the total intrinsic instance hardness:




Fig. 4. (a) Hardness is eradicated when the entire search space consists of solution paths: $\mathcal{G}_i = \mathcal{D}_i$. (b) Hardness arises as the deficit in the solution path diversity: $\mathcal{H}_i = \mathcal{G}_i - \mathcal{D}_i$. (c) Hardness is maximized when no solution paths exist: $\mathcal{H}_i = \mathcal{G}_i$. In all cases, the entire search space is taken into account when computing \mathcal{G}_i , i.e. there is no pruning of nogoods

Proposition 3 (\mathbf{H}^n). *The sum of the conditional intrinsic instance hardness $\Delta\mathcal{H}_i$ over all levels i of the search space is the total intrinsic instance hardness \mathbf{H}^n and is identical among all problem instances with the same number of variables n and the same number of solutions $|\mathcal{S}|$.*

$$\mathbf{H}^n(\mathcal{S}) = \mathbf{G}^n(\mathcal{S}) - \mathbf{D}^n(\mathcal{S}) = \sum_{i=1}^n \Delta\mathcal{H}_i = n - \log |\mathcal{S}|$$

This is especially intriguing. According to Proposition 3, all instances with the same number of variables n and the same number of solutions $|\mathcal{S}|$ have the same amount of total intrinsic hardness \mathbf{H}^n . However, bear in mind that the individual $\Delta\mathcal{H}_i$'s which make up \mathbf{H}^n are not necessarily identical for all such instances (only isomorphic instances have equal $\Delta\mathcal{H}_i$'s for all levels). What this essentially entails is that, among instances with equal n and $|\mathcal{S}|$, the equal total intrinsic hardness \mathbf{H}^n is expended at different levels of the search space. A relevant discussion in Sect. 4.2 will show that instances with a uniform distribution of solutions have low $\Delta\mathcal{H}_i$ for low values of i . On the other hand, instances with clustered solutions have a uniform distribution of the intrinsic hardness $\Delta\mathcal{H}_i$ among all levels of the search space.

Proposition 3 also serves to obviate the intimate connection between the total intrinsic instance hardness \mathbf{H}^n and the constrainedness κ of a single instance:

$$\kappa = 1 - \frac{\log |\mathcal{S}|}{n} = \frac{\mathbf{H}^n(\mathcal{S})}{n} \quad (2)$$

Under this light, κ can be understood as the *average instance hardness over all levels of the search space* or, alternatively, as average deficit in the amount of choice available for selecting a solution path, throughout the search space. This development provides well-founded justification for the establishment of κ as a hardness predictor and verifies the following claim from [2]:

Our definition of κ generalizes a number of parameters introduced in a variety of problem classes. This suggests that constrainedness is a fundamental property of problem ensembles.

Unfortunately, both \mathbf{H}^n and κ suffer from the same deficiency in characterizing instance hardness: they *fail to discriminate between instances with non-isomorphic solution sets*. This is due to the fact that \mathbf{H}^n and κ are, respectively, a sum and an average over the levels of the search space and thus fail to retain any information about the behavior of the individual $\Delta\mathcal{H}_i$'s. Recall that in Sect. 3.2 it was argued that the individual $\Delta\mathcal{H}_i$'s directly reflect the distribution of solutions and that they are necessary and sufficient for discriminating between non-isomorphic solution sets.

This provides an additional explanation as to the reason why there is significant variance in computational cost among problem instances belonging to the same ensemble: since κ depends only upon the number of solutions $|\mathcal{S}|$ and not their distribution, it is possible for non-isomorphic instances with different intrinsic instance hardness $\Delta\mathcal{H}_i$ per level to belong to the same ensemble.

The question then naturally arises: which hardness measure is not so crude as to be able to discriminate between instances with non-isomorphic solution sets? It can be argued, along the lines of [7], that a quantity such as the sum of \mathcal{H}_i 's can be employed in order to capture additional information about problem instances. In such a case, drawing from Sect. 4.2, the most difficult instances (among all instances with the same number of solutions $|\mathcal{S}|$) are the ones in which the solutions are *clustered*. This is in complete accordance with [11], [3] but further exploration is required. Note that *approximate entropy*, another information theoretic measure, is applied on problem structure in [12] in order to distinguish among ensemble instances with different distributions of solutions.

4.2 Bounds

The upper bound for solution path diversity \mathcal{D}_i corresponds to the case when the solution paths are uniformly distributed among the assignments of level i . The lower bound corresponds to the case when the solutions form a cluster of diameter c , thereby resulting in the greatest possible non-uniformity in the distribution of solution paths among the assignments of level i . The derivation of the lower bound is somewhat involved.

Proposition 4 (\mathcal{D}_i Bounds). *If $|\mathcal{S}|=2^c$, $0 \leq c \leq n$, then the following holds for solution path diversity \mathcal{D}_i at level i :*

$$\frac{c^i}{n} + \log C(n, i) \leq \mathcal{D}_i \leq \min\{i, c\} + \log C(n, i)$$

Using Proposition 4 and Defs. 5 and 6 it is trivial to compute bounds for the intrinsic instance hardness \mathcal{H}_i and $\Delta\mathcal{H}_i$.

Proposition 5 (\mathcal{H}_i , $\Delta\mathcal{H}_i$ Bounds). *If $|\mathcal{S}|=2^c$, $0 \leq c \leq n$, then the following holds for \mathcal{H}_i and $\Delta\mathcal{H}_i$ at level i :*

$$\max\{0, i - c\} \leq \mathcal{H}_i \leq \left(\frac{n - c}{n}\right) i$$

$$0 \leq \Delta\mathcal{H}_i \leq 1$$

Note that both bounds pertain exclusively to instances with non-empty solution sets and no pruning of nogoods. The case when the solution set is empty is trivial: intrinsic hardness \mathcal{H}_i coincides with the general path diversity \mathcal{G}_i . Illustrations are to be found in Fig. 5 for problem instances with 10 variables and 2^4 solutions.




Fig. 5. Intrinsic instance hardness \mathcal{H}_i bounds and conditional intrinsic instance hardness $\Delta\mathcal{H}_i$ bounds vs. search level i , for problem instances with 10 variables and 2^4 solutions. Additional values for \mathcal{H}_i and $\Delta\mathcal{H}_i$ computed for two specific instances are included

The bound for $\Delta\mathcal{H}_i$ reveals how informative a quantity it is and how it conveys information by mere inspection. Notice also how $\Delta\mathcal{H}_i$ remains constant at all levels in the case when the solutions are organized in a single cluster. Bearing in mind that $\Delta\mathcal{H}_i$ reflects an average over the instantiations of level i , this entails that the subproblems encountered throughout the search space are *persistently versatile*: there is ample margin for both error and success and algorithmic choices play the decisive part. On the other hand, when the solutions are as far apart as possible, there is such uniformity in the distribution of solution paths that the hardness $\Delta\mathcal{H}_i$ remains zero up to level c : algorithmic choices are inconsequential since all choices to be made are equally beneficial. However, the hardness $\Delta\mathcal{H}_i$ reaches its maximum beyond level c , with all choices being equally critical because of the spread and scarcity of solution paths.

Approximating Instance Hardness. The form of \mathcal{H}_i and $\Delta\mathcal{H}_i$ for specific instances (as depicted in Fig. 5) is especially intriguing. It becomes apparent that these quantities can be approximated by simple functions involving a control

parameter μ , which characterizes the uniformity of solution distribution in the set \mathcal{S} . This behavior is due to the direct association between the distribution of solution paths among consecutive levels and is also in accordance with very similar observations made in [7]. An initial attempt was to employ the sigmoid:

$$\Delta\mathcal{H}_i \approx \frac{1}{1 + \left(\frac{c}{n-c}\right)2^{-\mu(i-c)}}$$

For values $\mu \rightarrow 0$ (uniform solution distribution) and $\mu \rightarrow \infty$ (perfect clustering), the approximation coincides with the derived bounds. Unfortunately, for $0 < \mu < 1$, the approximation is unsatisfactory. This remains a stimulating open issue.

5 Conclusions

An insightful hint in [13] provides a condensed explanation of what this research has to offer:

In addition to the *number of solutions*, their clustering also contributes to the variance in the search cost. [...] It thus remains an open question whether additional parameters will be required to specify the clustering sufficiently tightly even if the number of solutions were well specified. [...] A better specification of the number and location of solutions requires more information about the structure of the problems, but is independent of the search method used. However, search methods themselves can differ in how well they avoid unproductive choices [...].

This is exactly the gap that is now bridged by intrinsic instance hardness. It is an intuitively appealing, algorithm-independent measure of instance hardness which directly reflects the clustering of solutions. At the same time, it justifies and supersedes constrainedness, which ignores the distribution of solutions and pertains only to their number. Novel issues are raised regarding constructive search and the constituents of hardness, especially with the bounds and rudimentary approximation results of section 4.2 in mind.

References

1. Williams, C., Hogg, T.: Exploiting the Deep Structure of Constraint Problems. *Artificial Intelligence* **70** (1994) 73–117
2. Gent, I.P., MacIntyre, E., Prosser, P., Walsh, T.: The Constrainedness of Search. In: *AAAI/IAAI*. Volume 1. (1996) 246–252
3. Parkes, A.J.: Clustering at the Phase Transition. In: *AAAI/IAAI*. (1997) 340–345
4. Shannon, C.E.: A Mathematical Theory of Communication. *The Bell Systems Technical Journal* **27** (1948) Reprinted with corrections
5. Slaney, J.: Is there a Constrainedness Knife-Edge? In: *Proceedings of the 14th European Conference on Artificial Intelligence*. (2000) 614–618

6. Hogg, T., Huberman, B., Williams, C.: Phase Transitions and the Search Problem. *Artificial Intelligence* **81** (1996) 1–15
7. Crutchfield, J., Feldman, D.: Regularities Unseen, Randomness Observed: Levels of Entropy Convergence. *Chaos* **13** (2003) 25–54
8. Vlasie, D.: The Very Particular Structure of the Very Hard Instances. In: AAAI/IAAI. Volume 1. (1996) 266–270
9. Ornstein, D.: Bernoulli Shifts with the Same Entropy are Isomorphic. *Adv. in Math.* **4** (1970) 337–352
10. Walsh, T.: The Constrainedness Knife-Edge. In: AAAI/IAAI. (1998) 406–411
11. Mezard, M., Ricci-Tersenghi, F., Zecchina, R.: Alternative Solutions to Diluted p-Spin Models and Xorsat Problems. *Journal of Statistical Physics* (2002)
12. Hogg, T.: Which Search Problems are Random? In: Proceedings of AAAI98. (1998)
13. Hogg, T.: Refining the Phase Transitions in Combinatorial Search. *Artificial Intelligence* **81** (1996) 127–154