Sudoku Players' Forums

THE REAL DISTRIBUTION OF MINIMAL PUZZLES

Goto page Previous 1, 2, 3 ... 21, 22, 23, 24 Next

newtopic postreply **Sudoku Players' Forums Forum Index -> General/puzzle**

View previous topic :: View next topic

Author	Message

Red Ed

Joined: 06 Jun 2005
Posts: 740

Posted: Fri Sep 25, 2009 9:45 am Post subject: quote

David P Bird wrote:

> If we know what the distribution of some characterisation measure should be for an entire population, we can see if the sample we're using is biased or not by checking if the sample distribution is in accordance with what it should be - a nice straightforward approach.

My yardstick for solution grids is number of patterns of certain types. I have no yardstick for puzzles.

Quote:

> While I'm doing my homework, perhaps you could turn your mind around about using flying start approaches - not what's wrong with them, but the possibilities of using one to take advantage of the time savings that could be had. It doesn't have to be the one I've suggested.

I'm pretty happy with the speed of my unbiased solution grids generator, which is I think a bit quicker than suexg's. I hope you won't chide me for declining your suggestion.

Back to top profile pm

Red Ed

Joined: 06 Jun 2005
Posts: 740

Posted: Fri Sep 25, 2009 10:16 am Post subject: quote

eleven, thanks for the code.

I commented out this section:

Code:

```
nClues=81;

part++;

for(i=1;i<=81;i++){mr4:x=(MWC>>8)&127;if(x>=i)goto
mr4;x++;P[i]=P[x];P[x]=i;}
for(i1=1;i1<=81;i1++){s1=A[P[i1]];if(s1){A[P[i1]]=0;
if(--nClues==34){if(solve()>1)goto m0;}if(nClues<
34&&solve()>1){A[P[i1]]=s1;break;}}}


i=++i1;for(i1=i;i1<=81;i1++){s1=A[P[i1]];if(s1){A[P[i1]]=0;if(solve()<2)goto
m0; A[P[i1]]=s1;}}
```

Having done that, bias in the complete grids so generated was as follows:

Code:

```
progress: 1000000 grids
  +-----------------------------+----------+-----------+----------+
  |           Pattern           |   Bias   |  Z-score  |   Z_1M   |
  +-----------------------------+----------+-----------+----------+
  | 1..2......6...19.......9..2 |   -5.87% |   -133.42 |  -133.42 |
  | .....2..8.............8..2 |   -9.77% |   -130.99 |  -130.99 |
  | 5....83........5...3....8.. |    7.17% |    130.08 |   130.08 |
  | 7..8..9...8..7....93.5....8 |   -5.67% |   -128.54 |  -128.54 |
  | 6.2..4....79..8.65.58..1..7 |  -10.14% |   -127.97 |  -127.97 |
  | .5........3.............35 |   -3.94% |   -127.78 |  -127.78 |
```

```
| .........6......4.4..6..... |    2.95% |    127.78 |   127.78 |
| 2...7..4..4......27........ |   -3.77% |   -127.66 |  -127.66 |
| .34.9.87..5......289..35.6. |   -9.35% |   -127.37 |  -127.37 |
| ...9.71...........71.....9. |   -2.89% |   -126.87 |  -126.87 |
| .87...1.22.........6124875. |   -9.30% |   -126.67 |  -126.67 |
| 318..79...4......77.59.3.2. |  -10.83% |   -126.61 |  -126.61 |
| ...7..4...7..4.9.5..9...... |   -4.64% |   -126.57 |  -126.57 |
| .1.7862...3.2.4.....7.91.58 |   -5.46% |   -126.32 |  -126.32 |
| 218.37...39..24..7.......8. |  -10.23% |   -126.17 |  -126.17 |
| 3.1..975..4..6...3.97.2.1.. |  -10.03% |   -126.01 |  -126.01 |
| 6...8..7...87.2.1.7..3...6. |  -10.48% |   -123.61 |  -123.61 |
| 578.613...3....764.2..9..1. |  -10.65% |   -122.80 |  -122.80 |
| ..2.9.1.....4....2..9.21... |   -8.70% |   -122.78 |  -122.78 |
| ..7....8....6.8....6.....7. |   -4.32% |   -122.71 |  -122.71 |
+-----------------------------+----------+-----------+----------+
```

I don't have the old results to hand, but I think these new results are just as bad as ever. Would be interesting to know if the same *type* of bias is showing through.

EDIT: no, actually these results are better than the old ones. Replacing prefix() with an empty function gives Z_1M = 200ish. And, by eye at least, the top however-many bias patterns appear to be different in the with/without prefix() cases.

EDIT2: gosh, what's more interesting is that the nature of the bias has switched. With prefix(), the greatest bias is in under-represented patterns; without prefix(), the greatest bias is in over-represented patterns. How strange! Here are the without-prefix() results:

Code:

```
progress: 1000000 grids
  +-----------------------------+----------+-----------+----------+
  |           Pattern           |   Bias   |  Z-score  |   Z_1M   |
  +-----------------------------+----------+-----------+----------+
  | .51.82...3.....82.97....51. |   25.22% |    208.44 |   208.44 |
  | 4.9....6......649.......... |   12.06% |    205.19 |   205.19 |
  | ...9.71...........71.....9. |   -4.66% |   -204.90 |  -204.90 |
  | 5....1..2.472.9..8..9...7.4 |   24.66% |    202.63 |   202.63 |
  | .4.8.567..672.43...18..9... |   18.65% |    201.13 |   201.13 |
  | .....6..2..17.89.69.....7.8 |   27.93% |    200.92 |   200.92 |
  | .1....9......9.1.7...7..... |   12.15% |    200.18 |   200.18 |
  | ..2.4..6.8.62.1..9479....5. |   18.68% |    197.36 |   197.36 |
  | ..9....1.....9...8....81... |   -4.80% |   -196.89 |  -196.89 |
  | ...3.295.713...6...8...61.7 |  102.22% |    196.69 |   196.69 |
  | ...934.7..3..1...6715.2...9 |   24.68% |    196.68 |   196.68 |
  | .35.....7....4.6.196.87.3.5 |   21.04% |    196.47 |   196.47 |
  | .8.6759....4...1.39...1.7.5 |   21.66% |    196.25 |   196.25 |
  | 13..9...42.51...9.7..25...8 |   20.96% |    195.78 |   195.78 |
  | 27....614...91...714...238. |   26.13% |    195.47 |   195.47 |
  | .49..................4.9.. |   -3.77% |   -195.22 |  -195.22 |
  | ....61.................16. |    6.78% |    195.22 |   195.22 |
  | ......476...64.....7....... |  101.09% |    195.22 |   195.22 |
  | 28..3.6......7....93.6..28. |   18.89% |    193.24 |   193.24 |
  | 2.3....19..7358....647..8.2 |  105.60% |    193.05 |   193.05 |
  +-----------------------------+----------+-----------+----------+
```

Back to top profile pm

eleven

Joined: 10 Feb 2008
Posts: 495

Posted: Fri Sep 25, 2009 11:11 am Post subject: quote

Exciting, so we could adjust the DLX solver with a good selection of randomly prefilled cells to give an optimum in the grid pattern distribution (which along the way makes the generator faster) ?

Back to top profile pm

Red Ed

Joined: 06 Jun 2005
Posts: 740

Posted: Fri Sep 25, 2009 11:30 am Post subject: quote

Crumbs ... I wonder?! Well I doubt you'll get the thing unbiased that way, but you may well be able to bring the bias down. I'll play with changing the number of prefilled cells, just out of curiosity.

Back to top profile pm

Red Ed

Posted: Fri Sep 25, 2009 11:48 am Post subject: quote

Joined: 06 Jun 2005
Posts: 740

Hmm. Odd. If you prefill only the first 20 of those 21 positions then the Z_1M bias score switches from -130 to +/-150, with a healthy(?) mix of +/- scores in the displayed list.

I find that impossible to understand at the moment **:-o**

EDIT: prefilling just the leading diagonal produces Z_1M = 110. Code:

Code:

```
void prefill()
{
int i,k,o[10];
for(i=1;i<=9;i++){do k=(MWC>>9)&15;while
(k>=i);k++;o[i]=o[k];o[k]=i;}
for (i=0;i<3;i++) A[1+10*i]=o[i+1];
for(i=1;i<=9;i++){do
k=(MWC>>9)&15;while(k>=i);k++;o[i]=o[k];o[k]=i;}
for (i=0;i<3;i++) A[31+10*i]=o[i+1];
for(i=1;i<=9;i++){do
k=(MWC>>9)&15;while(k>=i);k++;o[i]=o[k];o[k]=i;}
for (i=0;i<3;i++) A[61+10*i]=o[i+1];
}
```

Back to top profile pm

eleven Posted: Fri Sep 25, 2009 12:24 pm Post subject: quote

Joined: 10 Feb 2008
Posts: 495

Thanks, (because) of course i should have said "...*you* could adjust ..." - and i understand, thats it out of your interest.

To come back to an earlier question. **gsf** has a complete collection of grids. It should not be too much effort for him to randomly select a sample from it. How much grids would you need to check the bias ?

Back to top profile pm

Red Ed Posted: Fri Sep 25, 2009 12:33 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 740

Ah, now that's an interesting one. His grids are canonicalised. If you just picked one flat randomly (+ random relabelling & row/col etc. swaps) then you would still be biased, but only very, very, slightly. The bias in that case comes from the fact that a small percentage of grids are automorphic. I'm not sure how many grids would be required for my program to detect the bias: lots, probably!

But the real reason not to generate solution grids that way is that a lookup into the huge table that gsf has got would be more costly than using the B2347 algorithm mentioned earlier for unbiased generation.

Back to top profile pm

eleven Posted: Fri Sep 25, 2009 12:52 pm Post subject: quote

Joined: 10 Feb 2008
Posts: 495

I see, thanks again.

Back to top profile pm

David P Bird Posted: Fri Sep 25, 2009 3:33 pm Post subject: quote

Joined: 16 Sep 2008
Posts: 156
Location: Middle England

Red Ed, I'm stumped how to take our dialogue any further unless you are willing to educate me about how you select the patterns you employ in your bias estimates (which change from one run to another) and a lot more. The impression I've got is that you have neither the time nor the inclination to do that for me or indeed anyone else.

The point I made which you have either yet failed to understand, or have chosen to dismiss (maybe for a very valid reason, but unexplained) is this: If we can establish the frequencies that a number of fixed patterns or any other suitable characteristics in the solution grid for the whole population - as you have previously indicated you can do - we don't need to count possible completions for every trial to get an idea whether any particular solution grid generator shows any bias or not, we just

check how well their outputs match those frequencies.

You've said that while you are happy to spend time running trials that are relatively easy for you to set up - for which we should all thank you - you understandably aren't prepared to take a lot of time up coding something new. But I can't tell which side of your dividing line such trials would be!

The time I would need to spend to get familiar enough with this whole subject unaided is sadly beyond me, but I see **eleven** is covering some of the same ground anyway, so regretfully I'll bow out now.

Back to top profile pm

denis_berthier Posted: Fri Sep 25, 2009 8:59 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 842
Location: Paris, France

> **Red Ed wrote:**
>
> > **denis_berthier wrote:**
> >
> > > recall where my approach and Red Ed's differed (that is to say, before he finally adopted my controlled-bias approach).
> >
> > No, I have not adopted your approach. First, I don't aim to generate puzzles. Second, the technique I use for estimating the number-of-clues distribution was -- and remains (if I ever run it again) -- the c-clue-subsets-of-s-clue-subgrids method.

Which leads you to draw conclusions about the unbiased mean from a sample in which only **0.49 %** of the puzzles are above this mean instead of **21.5 %** in my approach. You can generate a sample one million times larger, you won't get much more than these 0.49 %.

> **Red Ed wrote:**
>
> btw, why are you scared of "astronomical" numbers?

Only in this case, because they lead you to the above 0.49 %.

> **Red Ed wrote:**
>
> Your puzzles have the additional nice property that they are *uncorrelated*.

Yes, this is the minimum one can ask from a random collection. But they first have the still nicer property of having a distribution not too far from the unbiased one (see figures above), thus allowing non-astronomical correction factors (large ones would entail unstability of the computations). After your old claims that the mathematics were the same [which they are if you understand "mathematics" only as using the obvious formula 81! / n! / (81-n)!], it is useful to recall this fact. This shows that the methods are indeed very different.

> **Red Ed wrote:**
>
> > **denis_berthier wrote:**
> >
> > > I've previously given several intuitive reasons why the controlled-bias algorithm can eliminate some bias in the complete grids (even large). This intuition has been confirmed (ironically, by Red Ed, while he was trying to prove the contrary)
> >
> > Why can't you grasp this? As I keep explaining, the only thing I was trying to prove was that the bias was real. I agree that the size (small) of the bias was a surprise to me, but in no sense does this mean that I failed to prove what I set out to prove.

The important part of my sentence was "This intuition has been confirmed". You comment only my comment. But you already answered yourself:

> **Red Ed wrote:**
>
> But, hey, maybe this is just nit-picking again.

Let me come back to an older post:

> **Red Ed wrote:**
>
> When you are quoting estimates from 10000 puzzles, I doubt that it makes any practical difference (for the proportions question) whether *suexg* or an unbiased solution grid source is used, because the variance due to the source will -- I think -- be small compared to variance due to the limited sample size. And since the latter type of variance is always going to be pretty high (because we're both pretty bad at finding 29- and 30-clue minimals), the upshot is that the bias in *suexg* is probably never really

> going to bite you.

10000 is generally considered as a very large sample - except for statistics of extremes.
My estimates are now from 250000 puzzles (using the accelerated version for the last 70000). And I'm not so bad at finding 29s and 30s. I have 1289 29s, 68 30s and even 1 31s.
If you have any doubt about your analysis due to my "limited" sample size, you can redo your comparisons with my updated distribution.

> **Red Ed wrote:**
>
> It'll bite me, because I want to count minimals, not just get their relative proportions; but it won't bite you.

Much ado about nothing. That's exactly what I was explaining.

Back to top profile pm www

Red Ed

Posted: Fri Sep 25, 2009 10:17 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 740

> **denis_berthier wrote:**
>
> ... leads you to draw conclusions about the unbiased mean from a sample in which only **0.49 %** of the puzzles are above this mean instead of **21.5 %** in my approach.

You're referring to my first post, the one that tipped you off that the existing mean estimates were too low and that spurred this whole thread. The c-clue-subsets-of-s-clue-subgrids method is a more focussed estimator for the parts of the distribution of most interest.

> **Quote:**
>
> > **Red Ed wrote:**
> >
> > Your puzzles have the additional nice property that they are *uncorrelated*.
>
> Yes, this is the minimum one can ask from a random collection.

When estimating mean values, you do **not** need the samples to be uncorrelated. It would be worth us both trying to understand whether paths (uncorrelated) or trees (correlated) works best for the right tail of the distribution.

> **Quote:**
>
> The important part of my sentence was "This intuition has been confirmed". You comment only my comment.

I was pegging you back on your means of address, which throughout the whole post, a propos of nothing, was to be confrontational.

> **Quote:**
>
> My estimates are now from 250000 puzzles (using the accelerated version for the last 70000). And I'm not so bad at finding 29s and 30s. I have 1289 29s, 68 30s and even 1 31s.

Can you quantify the confidence that you now have in your estimate of the mean number of clues?

Back to top profile pm

Red Ed

Posted: Fri Sep 25, 2009 10:23 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 740

> **David P Bird wrote:**
>
> **Red Ed**, I'm stumped how to take our dialogue any further unless you are willing to educate me about how you select the patterns you employ in your bias estimates (which change from one run to another) and a lot more. The impression I've got is that you have neither the time nor the inclination to do that for me or indeed anyone else.

If you want to keep asking questions for me to continue answering (as I've been doing quite diligently, I thought, despite your carping that I refuse to lavish you with enough attention) then take it to the Unbiased Grid Generation thread.

PS: re checking frequencies - yes, that's exactly what the bias tester does. Counting completions is what the generator does. Two different programs, two different methods, two different goals.

Back to top profile pm

denis_berthier

Posted: Fri Sep 25, 2009 10:55 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 842
Location: Paris, France

Red Ed wrote:

> **denis_berthier wrote:**
>
> > ... leads you to draw conclusions about the unbiased mean from a sample in which only **0.49 %** of the puzzles are above this mean instead of **21.5 %** in my approach.
>
> The c-clue-subsets-of-s-clue-subgrids method is a more focussed estimator for the parts of the distribution of most interest.

With the above figures, I don't see how that could be true.

Red Ed wrote:

> I was pegging you back on your means of address, which throughout the whole post, a propos of nothing, was to be confrontational.

Re-read your first posts . Or maybe you think saying that the methods are different and giving the above objective figures is confrontational? Well, it seems that both of us should try to be less confrontational (but it also seems we both enjoy it, don't we?).

Red Ed wrote:

> **denis_berthier wrote:**
>
> > My estimates are now from 250000 puzzles (using the accelerated version for the last 70000). And I'm not so bad at finding 29s and 30s. I have 1289 29s, 68 30s and even 1 31s.
>
> Can you quantify the confidence that you now have in your estimate of the mean number of clues?

No more than you can with your method.
I just observe that, as the sample grows, the estimated unbiased mean doesn't change significantly.

Back to top profile pm www

Red Ed

Posted: Fri Sep 25, 2009 11:20 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 740

Re c-clue subsets: it's a *different method* to the one in my first post that you referred to.

On this ...

Quote:

> **Red Ed wrote:**
>
> > Can you quantify the confidence that you now have in your estimate of the mean number of clues?
>
> No more than you can with your method.
> I just observe that, as the sample grows, the estimated unbiased mean doesn't change significantly.

... it would be good to quantify the confidence/uncertainty. Why not try bootstrap resampling.

Back to top profile pm

denis_berthier

Posted: Fri Sep 25, 2009 11:23 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 842
Location: Paris, France

UPDATED RESULTS FOR THE CONTROLLED-BIAS GENERATOR
suexg-cb was used to produce 250,000 minimal puzzles, the first 180,000 of which were obtained before the accelerated version (deleting the first 46 clues without doing any test) was available.

Here are the results for the number-of-clues distribution

raw-average = 25.65 unbiased-average = 26.56

Code:

```
#clues  raw-dist   unbiased-dist
19      0          0.0 (*)
20      2          0.03 (*)
21      3          0.14 (*)
22      323        40.51
23      4953       1594
```

24	30359	23604
25	77083	136646
26	83714	319632
27	41986	326555
28	10219	153284
29	1289	35336
30	68	3231
31	1	78 (*)
32	0	0 (*)

* values based on few data are not reliable.

PS.: The estimated mean SER and NRCZT are unchanged

Back to top profile pm www

Display posts from previous: All Posts Oldest First Go

newtopic postreply Sudoku Players' Forums Forum Index -> General/puzzle

All times are GMT - 8 Hours

Goto page Previous 1, 2, 3 ... 21, 22, 23, 24 Next

Page 22 of 24

Jump to: General/puzzle Go