Randomness and Computation 2018/19
Week 10 tutorial sheet (12-1pm, Tues 26th, Wed 27th March)

1. For this one, we need to examine the equations we used to analyse the time to reach a satisfying 2SAT assignment (when one exists). If you remember, in lecture 14, we derived a system of equations characterising "hitting times" (expected time to hit state n, thus satisfying all clauses) from state j, $0 \leq j \leq n-1$.

   We were able to show $h_n = 0$ and $h_0 = h_1 + 1$ for the end cases.

   For all the in-between states $j = 1, \ldots, n-1$, we had the relationship $h_j = \frac{1}{2}(h_{j+1} + h_{j-1} + 1)$. Further workings allowed us to show that

$$h_k = h_{k+1} + (2k+1),$$

   this also revealing that the largest hitting times were from the lowest-valued states k. We then went on to derive the bound on hitting time from 0.

   If our algorithm changes to start from a uniform random assignment, then we don't expect to start at $k = 0$. In fact, the expected number of *correctly assigned variables* from this assignment will be exactly $\frac{n}{2}$. Also we'll be able to show that the number of assignments with score j away from $n/2$ decreases exponentially as j increases. the probability that we get an assignment with *exactly* $n/2$ variables "correct" (wrt the hidden perfect assignment) is

$$\frac{n \cdot \ldots \cdot (n/2+1)}{(n/2)!2^n},$$

   whereas the probability of getting an assignment with only $n/4$ variables correct (with the same result if we ask for $3n/4$ variables correct) is much lower at

$$\frac{n \cdot \ldots \cdot (3n/4+1)}{(n/4)!2^n}.$$

   Suppose we consider for now, the expected hitting time from $k = n/2$. Then applying $h_k = h_{k+1} + (2k+1)$ repeatedly, we will obtain the result that

$$h_{n/2} = \sum_{k=n/2}^{n-1}(2k+1) = (n/2) + 2[\sum_{k=1}^{n-1} - \sum_{k=1}^{n/2}] = (n/2) + 2[\frac{n(n-1)}{2} - \frac{n(n/2+1)}{4}] = \frac{3n^2}{4}.$$

   Note that this is not a huge amount better than the result we got for starting at h_0 (that was n^2).

   We didn't do the analysis of the *weighted* value of the expected k (after having the assignment) but the majority of the probability distribution on $\{0, \ldots, n\}$ (after having drawn a random assignment and checked number of variables wrong) will be strongly concentrated around $n/2$, so this result is meaningful.

   Overall, the answer is "not a lot" (in answer to how this would affect the running time to a satisfying assignment).

2. Question about vertex colourings.

   (a) We have a 3-colourable graph (can have its vertices *properly coloured* with 3 colours) and we are asked to create a 2-colourings where there is no triangle ($\{u, v, w\}$ with all edges present) which is "all one colour".

   Ok, as suggested by the hint, we start with the given 3-colouring. Choose any of the three colours to eliminate - say "green" - and decide to re-colour all of the previously green vemrtices with "blue". This leaves a graph in which every veortex is "blue" or "red", so 2-coloured. Now we must prove that we have no monochromatic triangle. Well if there was some all-"red" triangle in the graph, this would have been all-"rerd" in the original 3-colouring (we did not add any extra red vertices), and very definitely would not have been a 3-colouring originally. So this isn't possible. Alternatively suppose we had triangle all-"blue" come about as a consequence of the re-coouring of some greens. Well, then that tells us that we only had "blue" and "green" in that triangle, before we coloured. However, the triangle has 3 vertices, and all are connected, so two colours could not have been enough to colour it. So again, even in the re-colouring, it can't be all-"blue".

   Hence no triangles are monochromatic in the re-colouring.

   (b) We are asked now to consider starting with an arbitrary 2-colouring of G (which is still a 3-colourable graph) which might have some monochromatic triangles. We are asked to consider the random process where we choose a monochromatic triangle (from whichever appear/remain in G) and flip the colour of any of the three vertices (each has $1/3$ chance of being flipped). Our goal is to bound the number of flips needed to achieve a 2-colouring where no triangle is monochrome.

   The analysis of this process should be viewed in a similar way to that of the 2SAT problem, in relation to some underlying 2-colouring which contains no monochromatic triangles (since we believe such a 2-colouring does exist). However, the details are harder/messier than for 2DNF. We start with an initial colouring $c : V \to \{\mathtt{blue}, \mathtt{red}\}$ and we consider the evolution of $c_0 = c, c_1, \ldots, c_t, \ldots$ until we have a 2-colouring with no monochromatic triangles.

   A key statistic of the current colouring will be Y_t, the number of vertices in c_t with colours different to the target. We are going to use the trick suggested in the problem statement - instead of measuring "similarity" between c_t and some hidden monochromatic 2-colouring, we will compare c_t to a hidden proper 3-colouring of G (which we know exists). Our measure of closeness will use the trick suggested in the hint; for every $v \in V$

   - If v has a "red" colour in our 3-coloured target, we will omit to check a match with the underlying red v colour, we just skip-over this vertex.
   - However, if v is "green"/"blue" in the hidden proper 3-colouring, then our measure Y_t adds 1 if the label assigned in c_t is the same, 0 otherwise.

   Note we have three vertices in every triangle, and one of these will definitely be "red" in the hidden 3-colouring, so this means we "ignore" one of the triangle's vertices in our scoring - this will actually help our analysis. Note also that when we achieve a 2-colouring which matches all "non-red" vertices of the proper 3-colouring, this 2-colouring is guaranteed to be without monochromatic triangles. So it is enough to show our process eventually satisfies a match in the modified scoring. For formality, for any integer value k that Y_t might take with the "check matches with all non-red vertices of the 3-colouring" measure of *similarity*, we will consider the random variable Z_k, this being the "number of steps" needed to reach the target 3-colouring (ignoring red vertices) from a colouring where Y_t equals k.

We define $h_k = \mathrm{E}[Z_k]$ for $k = 0, \ldots, all$ where all is the maximum value Y_t can take (the number of non-red vertices of the hidden 3-colouring). We know $all < n$ but the specific value will depend on the graph and the hidden 3-colouring.

Now we consider the options when we choose a monochromatic triangle $\{u, v, w\}$ in our current colouring c_t, and randomly flip one of the colours. We know that in our current colouring c_t, all of $\{u, v, w\}$ are "blue", or alternatively all of $\{u, v, w\}$ are "green". Let's assume all are currently "blue" without loss of generality. We also know that in our target 3-colouring, each of $\{u, v, w\}$ has a different colour, one red, one green, one blue. With probability $1/3$ we will flip the vertex whose hidden colour is "red", this does not change Y_t or Z_t at all ($Y_{t+1} = Y_t$ and the difference is still k). Alternatively, with probability $1/3$ we will flip the vertex with hidden colour "green", and this will create an extra match (as we have done "blue" $\rightarrow$ "green") so we move from Z_k to Z_{k+1} in this case. In the other case, we end up flipping a blue vertex which should be blue, so with probability $1/3$ we move from Z_k to Z_{k-1}. Hence we can follow the structure of our 2SAT argument and write

$$\mathrm{E}[Z_j] \;= \frac{1}{3}(\mathrm{E}[Z_{j+1}] + 1) + \frac{1}{3}(\mathrm{E}[Z_{j-1}] + 1) + \frac{1}{3}(\mathrm{E}[Z_j] + 1),$$

for $j = 1, \ldots, all - 1$.

We have omitted $j = all$ from our list of indices as when $j = all$, then we have no mono-chromatic triangle and $Z_{all} = 0$.

However, by the structure of the "monochromaticity" condition, if *all* non-red vertices of the graph *clash* with the hidden 3-colouring, then we *also* are in the situation where all triangle are non-monochromatic. Hence $Z_0 = 0$ also. This will mean that the structure of our system of equations is slightly different to that for the 2DNF (apart from constants being different).

This allows us to write the following system:

$$h_j \;= \begin{cases} 0 & j = all \text{ or } j = 0 \\ \frac{1}{2}(h_{j-1} + h_{j+1}) + \frac{3}{2} & \text{for } j = 1, \ldots, all - 1 \end{cases} \tag{1}$$

We can now solve for solutions to the linear system in two directions, first working from index 0 and upwards, then working from index $all - 1$ and downwards. We will prove two relationships. First we will show that for $i = 1, \ldots, all - 1$, we have:

$$h_i \;=\; \tfrac{i}{i+1} h_{i+1} + i \cdot \tfrac{3}{2}.$$

Then we will show that for all $j = all - 1, \ldots, 1$, that

$$h_j \;=\; \frac{all-j}{all-(j-1)} h_{j-1} + (all - j) \cdot \frac{3}{2}.$$

As an example, consider the system (2). For the base case $i = 1$, note that equation (1) tells us that $h_1 = \frac{1}{2}(h_0 + h_2) + \frac{3}{2}$. By $h_0 = 0$, this is equal to $\frac{h_2}{2} + 1 \cdot \frac{3}{2}$, as required.
Then for the inductive step, assume that the claim holds for some $k - 1 < all - 1$ and consider $i = k$. By equation (1) we have $h_k = \frac{h_{k-1}+h_{k+1}}{2} + \frac{3}{2}$. Now using the (IH) for $i = k - 1$, we can rewrite this as

$$h_k \;=\; \frac{k-1}{2k} h_k + (k-1) \cdot \frac{3}{4} + \frac{h_{k+1}}{2} + \frac{3}{2}$$

This is equivalent to

$$\frac{k+1}{2k} h_k \;=\; (k-1) \cdot \frac{3}{4} + \frac{h_{k+1}}{2} + \frac{3}{2}$$

and multiplying across and simplifying, equivalent to

$$\begin{aligned} h_k &= \frac{2k(k-1)}{(k+1)} \cdot \frac{3}{4} + \frac{kh_{k+1}}{(k+1)} + \frac{3k}{k+1} \\ &= \frac{k}{(k+1)} h_{k+1} + \frac{k}{(k+1)}((k-1) \cdot \frac{3}{2} + 3) \\ &= \frac{k}{(k+1)} h_{k+1} + k\frac{3}{2}, \end{aligned}$$

which is equation (2) for $i = k$. Hence by induction we have (2) for all the specified indices. The proof of the system (3) is similar.
Now consider any specific indices $m, m + 1$ that are neither 0 nor all. With $i = m$ equation (2) gives

$$h_m \;=\; \frac{m}{m+1} h_{m+1} + m \cdot \frac{3}{2}.$$

With $j = m + 1$, equation (3) gives

$$h_{m+1} \;=\; \frac{all-(m+1)}{all-m} h_m + (all - m) \cdot \frac{3}{2}.$$

We can multiply this by $\frac{m}{m+1}$ to have

$$\frac{m}{m+1} h_{m+1} \;=\; \frac{m(all-(m+1))}{(m+1)(all-m)} h_m + (all - m) \cdot \frac{m}{m+1} \frac{3}{2}.$$

Then substituting into our expression for h_m above we get

$$h_m \;=\; \frac{m(all-(m+1))}{(m+1)(all-m)} h_m + (all - m) \cdot \frac{m}{m+1} \frac{3}{2} + m \cdot \frac{3}{2}.$$

We next subtract the rhs h_m from the left to get

$$\left(1 - \frac{m(all-(m+1))}{(m+1)(all-m)}\right) h_m \;=\; (all - m) \cdot \frac{m}{m+1} \frac{3}{2} + m \cdot \frac{3}{2},$$

which simplifies to

$$\frac{\mathsf{all}}{(m+1)(\mathsf{all}-m)}h_m \quad = \quad (\mathsf{all}-m)\cdot\tfrac{m}{m+1}\tfrac{3}{2} + m\cdot\tfrac{3}{2},$$

and multiplying across by $(\mathsf{all}-m)(m+1)$,

$$\begin{aligned}
\mathsf{all}\cdot h_m &= (\mathsf{all}-m)(m+1)\cdot\left((\mathsf{all}-m)\cdot\tfrac{m}{m+1}\tfrac{3}{2} + m\cdot\tfrac{3}{2}\right)\\
&= \tfrac{3}{2}m\left((\mathsf{all}-m)^2 + (m+1)(\mathsf{all}-m)\right)\\
&= \tfrac{3}{2}m(\mathsf{all}-m)(\mathsf{all}+1).
\end{aligned}$$

Hence $h_m = \frac{m(\mathsf{all}-m)(\mathsf{all}+1)}{\mathsf{all}}\frac{3}{2} \leq (\mathsf{all}+1)(m+1)\frac{3}{2}$ for all m lying between 2 and $(\mathsf{all}-2)$. This is $\Theta(\mathsf{all}^2) = \Theta(n^2)$ (note the worst case is when m is about midway between 0 and all).

note: I know I've skipped over the $m = 1, m = \mathsf{all}-1$ corner cases but equation (2) or equation (3) can be applied to these to get the answer in terms of $m = 2, m = \mathsf{all}-2$.

3. (a) If we have an input $a_1, \ldots, a_n; b$, let us note that $\frac{|\Omega|}{2^n}$ is the probability that a single random trial (where the x_i are chosen uar from $\{0,1\}$) belongs to the set of knapsack solutions. If Ω's cardinality (which depends on the specific values of $a_1, \ldots, a_n$ in relation to b) is too small in relation to 2^n, then even taking a polynomial number (eg, n, or n^2, or $6n^4 \ldots$) of random trials, we will probably never generate an element inside Ω (a vector $\bar{x} \in \{0,1\}^n$ which satisfies the knapsack condition wrt $a_1, \ldots, a_n; b$). And hence the scaled estimate will return 0, which is a very poor estimate for $|\Omega|$.

   To be specific, consider the case where $b = \sqrt{b}$ and all a_i are exactly 1. Then the probability that a random sample from $\{0,1\}^n$ is a knapsack solution (ie, has no more than $\sqrt{n}$ 1s) is at most

$$\binom{n}{\lfloor\sqrt{n}\rfloor}\cdot\frac{1}{2^{n-\sqrt{n}}} \;\leq\; \frac{n\ldots(n-\lfloor\sqrt{n}\rfloor)}{\lfloor\sqrt{n}\rfloor!}\cdot\frac{1}{2^{n-\sqrt{n}}} \;\leq\; n^{\sqrt{n}}\cdot\frac{1}{2^{n-\sqrt{n}}},$$

   and by $n = 2^{\lg(n)}$, we know

$$n^{\sqrt{n}}\cdot\frac{1}{2^{n-\sqrt{n}}} \;=\; \frac{2^{\lg(n)\cdot\sqrt{n}}}{2^{n-\sqrt{n}}} \;=\; \left(\frac{2^{\lg(n)}}{2^{\sqrt{n}-1}}\right)^{\sqrt{n}}.$$

   However, for $n \geq 16$, we have $\lg(n) \leq \sqrt{n}$, and for n "sufficiently large" ($n \geq 512$ will do), we have $\sqrt{n}/2 - 1 \geq \lg(n)$, and then for this "$n$ sufficiently large" the probability above will be smaller than $(2^{-\sqrt{n}/2})^{\sqrt{n}}$, ie, smaller than $2^{-n/2}$.

   Clearly, we could never detect a success of probability $\leq 2^{-n/2}$ by taking polynomially-many samples.

   (b) Now we consider the Markov chain.

i. To show that the chain is *irreducible*, we need to show that for every pair of two feasible knapsack solutions $\bar{x}, \bar{z}$, there is a sequence of intermediate knapsack solutions $\bar{y}(0) = \bar{x}, \bar{y}(1), \dots, \bar{y}(\ell) = \bar{z}$ such that $M[\bar{y}(i), \bar{y}(i+1)] > 0$ for every $i, 0 \leq i \leq \ell - 1$. This is simpler to show for knapsack solutions than it was for contingency tables: define $D =_{\mathsf{def}} \{i : x_i = 1 \text{ and } z_i = 0\}$, and define $A =_{\mathsf{def}} \{i : x_i = 0 \text{ and } z_i = 1\}$.
Now order the indices in D as $i_1, \dots, i_{|D|}$ (in arbitrary fashion) and for $j = 1, \dots, |D|$, define $\bar{y}(j)$ to be the knapsack solution $\bar{y}(j-1)$ with the i_jth index now switched from 1 to 0.
Note that by definition of $\bar{y}(j+1)$, it contains one less 1 than $\bar{y}(j)$, and hence is a feasible solution if $\bar{y}(j)$ was feasible. We started at $\bar{y}(0) =_{\mathsf{def}} \bar{x}$ which was feasible, hence each of $\bar{y}(1), \dots, \bar{y}(|D|)$ will be feasible knapsack solutions belonging to Ω (as with each transition we have a knapsack with one fewer item). Also, by definition, each $\bar{y}(j)$ and $\bar{y}(j+1)$ differ by a single index i_j and hence are connected by a transition satisfying $M[\bar{y}(j), \bar{y}(j+1)] > 0$.
By definition, D is the set of indices which were set to 1 in $\bar{x}$ but missing in $\bar{z}$. Hence by definition $\bar{y}(|D|)$ is the knapsack solution which contains a 1-index in the position which are 1 in *both* $\bar{x}$ and $\bar{z}$. Hence $\bar{y}(|D|)$ is essentially the "intersection" of $\bar{x}$ and $\bar{z}$.
Now order the indices in A as $k_1, \dots, k_{|A|}$ (in arbitrary fashion) and for $h = 1, \dots, |A|$, define $\bar{y}(|D| + h)$ to be the knapsack solution $\bar{y}(|D| + h - 1)$ with the i_h-th index now switched from 0 to 1.
Note that by definition of the $\bar{y}(|D| + h)$, each such vector contains a 1 for each of the "intersection" indices plus also the indices $j_{|D|+1}, \dots, j_{|D|+h}$ which belonged to $\bar{z}$ but not to $\bar{x}$. Note this implies that the 1-indices of each $\bar{y}(|D| + h)$ are a subset of $\bar{z}$'s 1s, and hence each $\bar{y}(|D| + h)$ is a feasible solution. Also, by definition, each $\bar{y}(|D| + h)$ and $\bar{y}(|D| + h + 1)$ differ by a single index i_h and hence are connected by a transition satisfying $M[\bar{y}(|D| + h), \bar{y}(|D| + h + 1)] > 0$.
By construction we have shown there is a path between $\bar{x}$ and $\bar{z}$, for arbitrary feasible solutions $\bar{x}, \bar{z}$.
(this would have been neater had I phrased the knapsack solutions as "subsets of [n]" rather than in the 0/1 vector form, as I could have used set operations to define the details. Sorry about that)

ii. We now need to show that the chain is aperiodic.
For every pair of solutions $\bar{x}, \bar{z}$ we have a connecting path of length exactly $|D| + |A|$, as shown in i.. Also we know that at any step of this path which is not the all-1s state, that we can add two transitions by "adding index i, dropping index i" for a non-zero index. Hence a connecting path of length $|D| + |A| + 2$ is possible (it is also possible to addd/delete a few items to add $4, 6, 8, \dots$ to the length of the path). The gcd of path lengths is at most 2 (and if $|A| + |D|$ was odd, it is 1).
This is still not enough to be aperiodic if $|D| + |A|$ was even. We need one extra condition to show aperiodicity - that $\sum_{i=1}^{n} a_i > b$ (all items will not fit in the knapsack). With this assumption, we can take any solution along the path, and attempt "add index i" in succession for *all* indices which are initially 0 - at some point we will reach full capacity and the attempted addition will fail (since we

cannot fit all n items). On this last attempt we will have the same state/solution appear again without any change (in other words, at this stage there is a positive probability of transitioning to the same state). We can then remove all the extra items one-by-one to return to our original point on the original $(\bar{x}, \bar{z})$-path, but this time we have added an odd number of steps to the path to obtain a $(\bar{x}, \bar{z})$-path of odd length overall, and hence the gcd of all such paths must be 1.

iii. Since the chain has been shown to be ergodic, we know the stationary distribution must be unique. We consider $\pi(\cdot)$ where $\pi(\bar{x}) = \frac{1}{|\Omega|}$ for every feasible knapsack $\bar{x}$. If this is to be the stationary distribution, it must be that case that $\pi \cdot M = \pi$, in other words we must have

$$\pi(\bar{x}) = \sum_{\bar{y} \in \Omega} \pi(\bar{y}) M[\bar{y}, \bar{x}]$$

for every $\bar{x} \in \Omega$. Since $\pi(\cdot)$ is uniform, this will hold if and only if

$$\sum_{\bar{y} \in \Omega} M[\bar{y}, \bar{x}] = 1.$$

However, note that for any $\bar{x}, \bar{y}$ which are connected by a non-zero transition are both *feasible*, and are different in at most one index. Because both are feasible, we know that $M[\bar{x}, \bar{y}] = M[\bar{y}, \bar{x}] = \frac{1}{n}$. Hence

$$\sum_{\bar{y} \in \Omega} M[\bar{y}, \bar{x}] \; = \; \sum_{\bar{y} \in \Omega} M[\bar{x}, \bar{y}],$$

and this latter sum is equal to 1 by the properties of a transition matrix (all rows sum to 1).

Mary Cryan, polished 19th April