ECE 341 - Homework #14

Chi-Squared Tests. Summer 2024

Loaded Dice

1) The following Matlab code generates 180 random numbers from 1..6 (sixty 6-sided dice)

```
RESULT = zeros(1,6);
for i=1:180
   d6 = ceil(rand*6);
   RESULT(d6) = RESULT(d6) + 1;
end
```

Use a chi-squared test to determine if this is a fair die.

```
RESULT =     25     25     31     39     33     27
```

Computing the chi-squared score:

Roll	p	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
1	1/6	30	25	0.8333
2	1/6	30	25	0.8333
3	1/6	30	31	0.0333
4	1/6	30	39	2.7
5	1/6	30	33	0.3
6	1/6	30	27	0.3
			Total	5

Convert this to a probability using StatTrek (chi-squared table)

p = 0.58412

There is a 58.412% chance this is not a fair die (no conclusion)

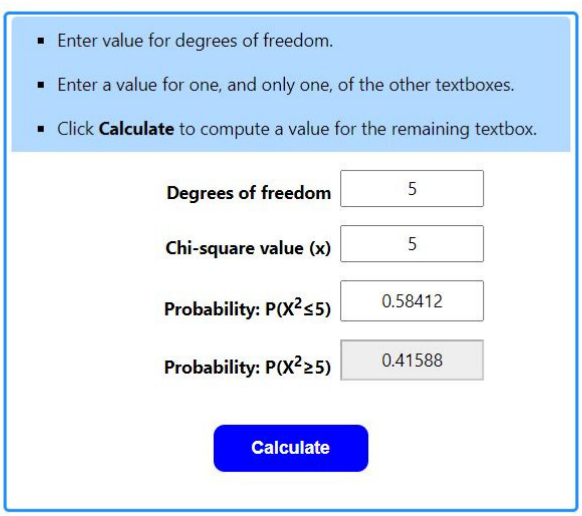

2) The following Matlab code generates 180 random die rolls with 8% loading (8% of the time you roll a 6)

```
RESULT = zeros(1,6);
for i=1:180
   if(rand < 0.08) d6 = 6;
   else d6 = ceil(rand*6);
   end
   RESULT(d6) = RESULT(d6) + 1;
end
```

Use a chi-squared test to determine if this is a fair die.

```
RESULT =   36    29    25    29    37    24
```

Roll	p	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
1	1/6	30	36	1.2
2	1/6	30	29	0.0333
3	1/6	30	25	0.8333
4	1/6	30	29	0.0333
5	1/6	30	37	1.6333
6	1/6	30	24	1.2
			Total	4.9333

From StatTrek, this corresponds to a probability of 0.58593

There is a 58.593% chance this die is loaded (no conclusion)

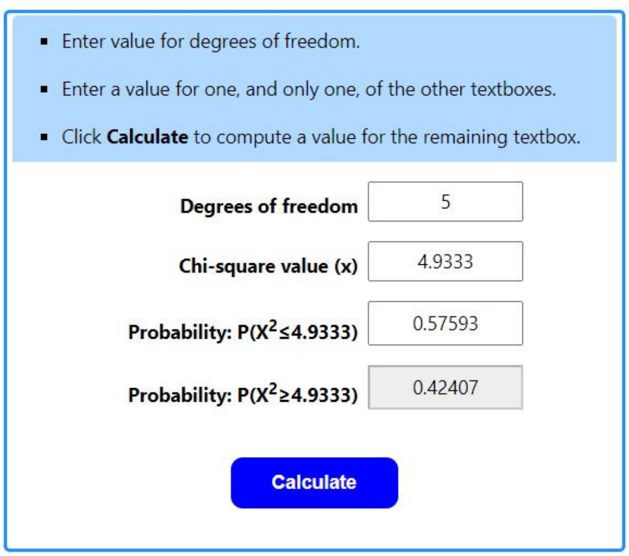

3) Repeat problem #2 with 1000 die rolls

```
RESULT =   163   145   154   151   156   231
```

Roll	p	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
1	1/6	166.67	163	0.0808
2	1/6	166.67	145	2.8175
3	1/6	166.67	154	0.9632
4	1/6	166.67	151	1.4733
5	1/6	166.67	156	0.6831
6	1/6	166.67	231	24.8296
			Total	30.8474

Now the probability is 0.99999

I am 99.999% certain this is a loaded die

With enough data, you can spot loaded dice

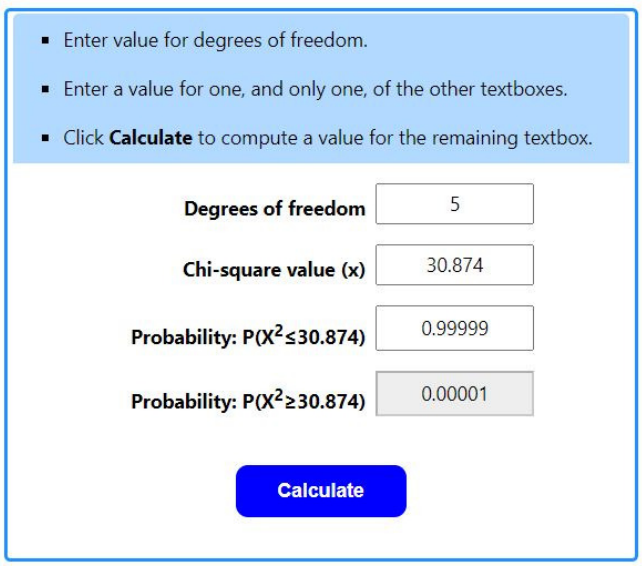

Am I psychic?

4) Shuffle a deck of 52 playing cards. Without looking at the top card, predict the suit (clubs, diamonds, hearts, and spades). Repeat for all 52 cards, keeping track of how many you got right and how many you got wrong.

- From the results, use a chi-squared test to determine if you are just guessing (25% chance of getting the suit correct.)

Result:

Correct: 15 times

Incorrect: 37 times

Compute the chi-squared score:

Result	p	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
Correct	1/4	13	15	0.3077
Incorrect	3/4	39	37	0.1026
			Total	0.4103

This corresponds to a probability of 0.47818

There is a 47.818% chance I'm not just guessing (no conclsion)

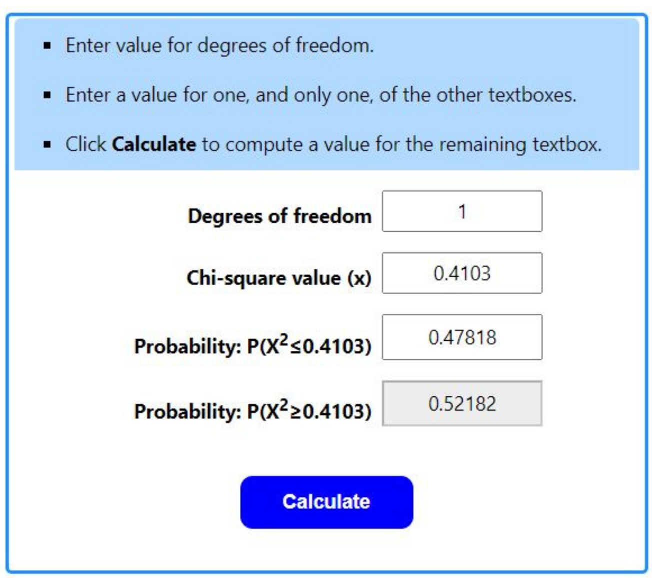

Central Limit Theorem:

5) The following code sums four uniform distributions

Y = sum(rand(4,1));

The Central Limit Theorem states that this will converge to a normal distribution with

- mean = 2.0
- variance = 4/12

Use a chi-squared test to determine if Y does / does not have a normal distribution.

Step 1: Collect data

- collected 100 values for Y

Step 2: Split the range space into N bins

- This is somewhat arbitrary
- I'll make each bin one standard deviation (1/2)

Step 3: Count how many times the data falls into each bin

Matlab Code

```
RESULT = zeros(1,4);
for n=1:100
    Y = sum(rand(4,1));
    bin = ceil(Y);
    RESULT(bin) = RESULT(bin) + 1;
end
RESULT =  1    48    47     4     0     0
```

Calculate the chi-squared score

bin (Y)	p	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
(0,1)	0.04165	4.165	1	2.4051
(1,2)	0.4584	45.84	48	0.1018
(2,3)	0.4584	45.84	47	0.0294
(3,4)	0.04165	4.165	4	0.0065
			Total	2.5428

From StatTrek, this corresponds to a probability of 0.53222

There isa 53.222% chance this is not a normal distribution (no conclusion)

Poisson approximation for a binomial distribution.

6) Let X be the number of 1's you get when you roll 60 dice. The Poisson approximation for the pdf is

$$\binom{60}{x}\left(\frac{1}{6}\right)^{x}\left(\frac{5}{6}\right)^{60-x} \approx \left(\frac{1}{x!}\right) 10^{x} e^{-10}$$

Use Matlab to count the number of 1's you get when you roll 60 dice

Repeat 200 times

- Check whether the result is consistent with a Poisson distribution with $\lambda = Np = 10$ using a Chi-squred test

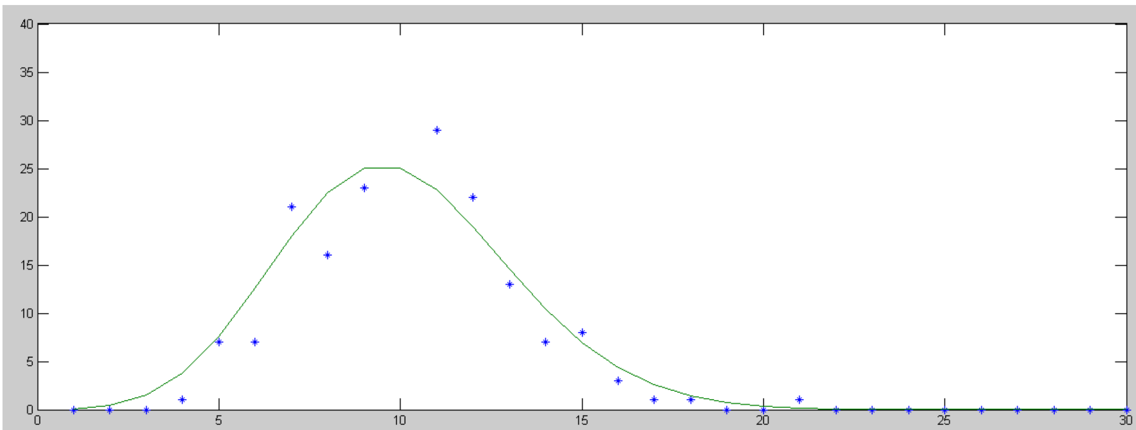

bin (Y)	n*p	N	$\chi^2 = \left(\frac{(np-N)^2}{np}\right)$
[0,4]	5.8415	1	4.0127
[5,9]	85.7354	74	1.6063
[10,14]	91.7224	111	4.0516
[15,19]	16.008	13	0.5652
[20,24]	0.6815	1	0.1489
[25,30]	0.0094	0	0.0094
		Total	10.3941

From StatTrek, this corresponds to a probability of 0.93519

There is a 93.519% chance the distributions are different