Mind map outline

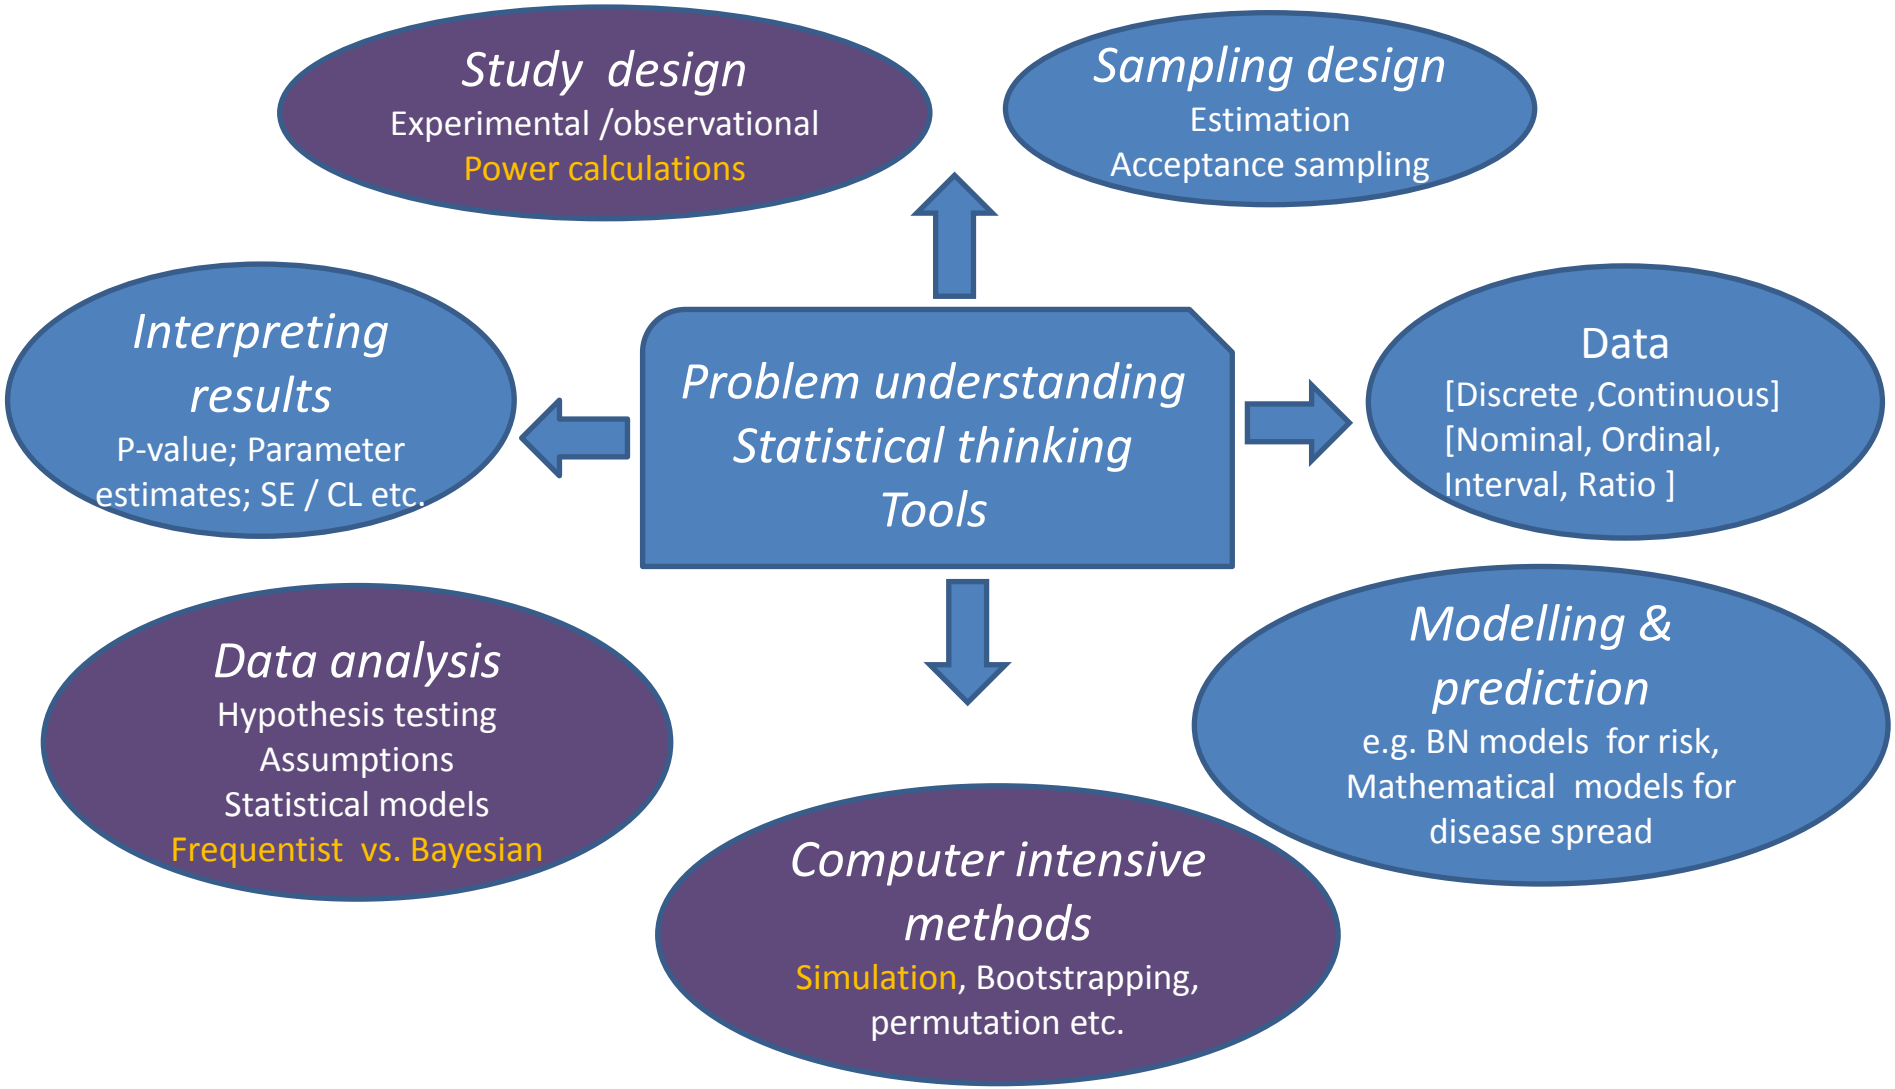

Power Analysis

What is Statistical Power?

• Statistical Power is the probability of rejecting the Null Hypothesis if it is false

e.g. The probability of getting a statistically significant treatment effect in your experiment if a true treatment effect exists.

•Formally Power= 1-β where β is the type II error.

•It depends on a number of factors: α –level, magnitude of the variation from the relevant sources, sample size, experimental design, and statistical test used.

Power Analysis

Why is power important?

- A prospective power analysis helps to decide the appropriate sample size for a trial.
- For many clinical trials a power of 80% or 90% is often required for ethical approval
- Retrospective or Post-hoc power analyses are less useful (and more controversial).

Power Analysis

How to calculate power

- Analytically

  e.g. For the difference in proportions with equal sized samples and α=0.05:

$$Z_{1-\beta} = \sqrt{\frac{n(p_1 - p_2)}{2\bar{p}(1-\bar{p})}} - 1.96$$

  – Most common statistical software packages include a power calculation module. There are also specialised programs available for the more complicated designs e.g. G*Power, Russ Lenth's Power and Sample size Java applet.

Power Analysis

How to calculate power

- Through simulation e.g.

```
res.prop.test<-rep(0,1000)
n<-110
for(i in 1:1000){
p1<-rbinom(1,n,.4)
p2<-rbinom(1,n,.6)
res.prop.test[i]<-prop.test(c(p1,p2),c(n,n))$p.value
}
sum(res.prop.test<0.05)/1000
[1] 0.818
```

(Analytical power=0.813)

- Can be very useful when the design or model is non standard and there is no analytical method easily available.

Power Analysis

Factors that influence power:

- The lower the α-level the lower the power.
- Power increases with sample size.
- The larger the minimum detectable difference the greater the power.
- The experimental design and model can influence the power e.g. Paired or two sample t-test

Power ↑ as the difference to be detected ↑

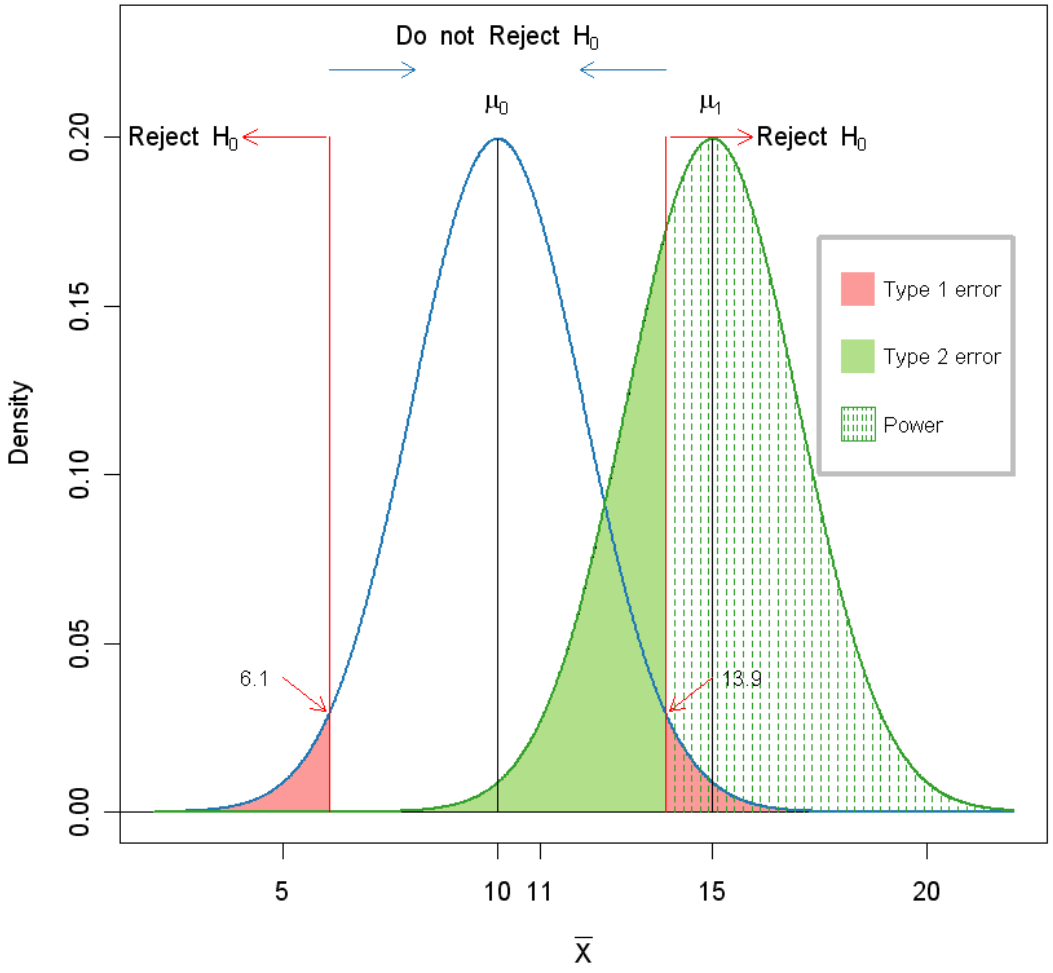

Power Analysis

http://homepage.stat.uiowa.edu/~rlenth/Power/

Java applets for power and sample size

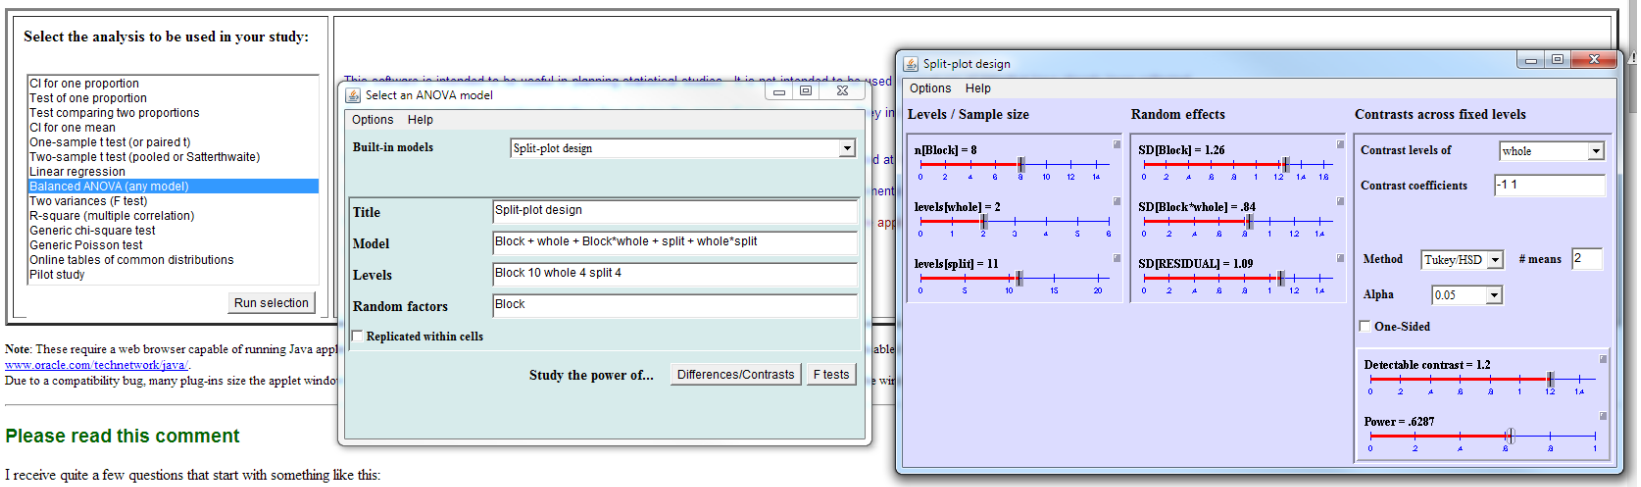

Note: These require a web browser capable of running Java appl
www.oracle.com/technetwork/java/.
Due to a compatibility bug, many plug-ins size the applet windo

Please read this comment

I receive quite a few questions that start with something like this:

"I'm not much of a stats person, but I tried [details...] -- am I doing it right?"

Please compare this with:

"I don't know much about heart surgery, but my wife is suffering from ... and I plan to operate ... can you advise me?"

Folks, just because you can plug numbers into a program doesn't change the fact that if you don't know what you're doing, you're almost guaranteed to get meaningless results -- if not dangerously misleading ones. Statistics really is like rocket science; it isn't easy, even to us who have studied it for a long time. Anybody who think it's easy surely lacks a deep enough knowledge to understand why it isn't! If your scientific integrity matters, and statistics is a mystery to you, then you need expert help. Find a statistician in your company or at a nearby university, and talk to her face-to-face if possible. It may well cost money. It's worth it.

If you're blocked by a security setting

You may get into a situation, especially after updating Java, where this applet is blocked by a security setting in your system. In that case, in order to run the applets, you need to this site to the list of exceptions to your security rules. To do this in Windows, find "Configure Java" on the Start menu, go to the security tab, and near the bottom, click on the "Edit site list" button, and add "http://homepage.stat.uiowa.edu" to the list. Then things should work.

Discussion group

Bayesian Statistics

- Frequentist hypothesis tests use the p-value as evidence against the Null hypothesis. It does not always mean that the alternative hypothesis is more likely.

An example: Suppose we have a test for PSA-X with sensitivity of 100% and specificity 99.9%. If we take a single sample at random and get a positive test result, then what is the probability that PSA-X is present?

Bayesian Statistics

- Frequentist hypothesis tests use the p-value as evidence against the Null hypothesis. It does not always mean that the alternative hypothesis is more likely.

An example: Suppose we have a test for PSA-X with sensitivity of 100% and specificity 99.9%. If we take a single sample at random and get a positive test result, then what is the probability that PSA-X is present?

Answer: We need more information! We need to know the probability of PSA-X occuring.

- If Pr(PSA-X)=0.5 then

$$\Pr\left(PSAX \mid T^{+}\right)=\frac{\Pr\left(T^{+} \mid PSAX\right)\times \Pr\left(PSAX\right)}{\Pr(T^{+})}=\frac{1\times 0.5}{1\times 0.5+0.001\times 0.5}=0.999$$

Bayesian Statistics

- Frequentist hypothesis tests use the p-value as evidence against the Null hypothesis. It does not always mean that the alternative hypothesis is more likely.

An example: Suppose we have a test for PSA-X with sensitivity of 100% and specificity 99.9%. If we take a single sample at random and get a positive test result, then what is the probability that PSA-X is present?

Answer: We need more information! We need to know the probability of PSA-X occuring.

- If Pr(PSA-X)=0.5 then

$$\Pr(PSAX \mid T^+) = \frac{\Pr(T^+ \mid PSAX) \times \Pr(PSAX)}{\Pr(T^+)} = \frac{1 \times 0.5}{1 \times 0.5 + 0.001 \times 0.5} = 0.999$$

- However if Pr(PSA-X)=1/1001 then

$$\Pr(PSAX \mid T^+) = \frac{1 \times \frac{1}{1001}}{1 \times \frac{1}{1001} + 0.001 \times \frac{1000}{1001}} = \frac{\frac{1}{1001}}{\frac{1}{1001} + \frac{1}{1001}} = 0.5$$

Bayesian Statistics

•Frequentists base their findings on the data and do not incorporate prior knowledge.

•In the previous example the probability of a positive test if PSA-X is not present is 0.1% and so a standard Hypothesis test would reject the Null.

•Bayesian statistics aims to use prior knowledge in addition to the data to make inferences about the unknown parameters using Bayes' Theorem

$$\Pr(\theta \mid data) = \frac{\Pr(data \mid \theta) \times \Pr(\theta)}{\int \Pr(data \mid \theta) \times \Pr(\theta) d\theta}$$

•Often the integral is intractable and so MCMC methods are used to sample from the posterior distribution using the relationship

$$Posterior \propto likelihood \times prior$$

Bayesian Statistics

•Once the model is set up a large number of observations are sampled from the posterior distribution.

•Since the earlier samples may be influenced by the starting values it is common practice to ignore the first batch (burn-in) of initial values.

•Base inferences on the posterior sample e.g. Median, mean, credible intervals...

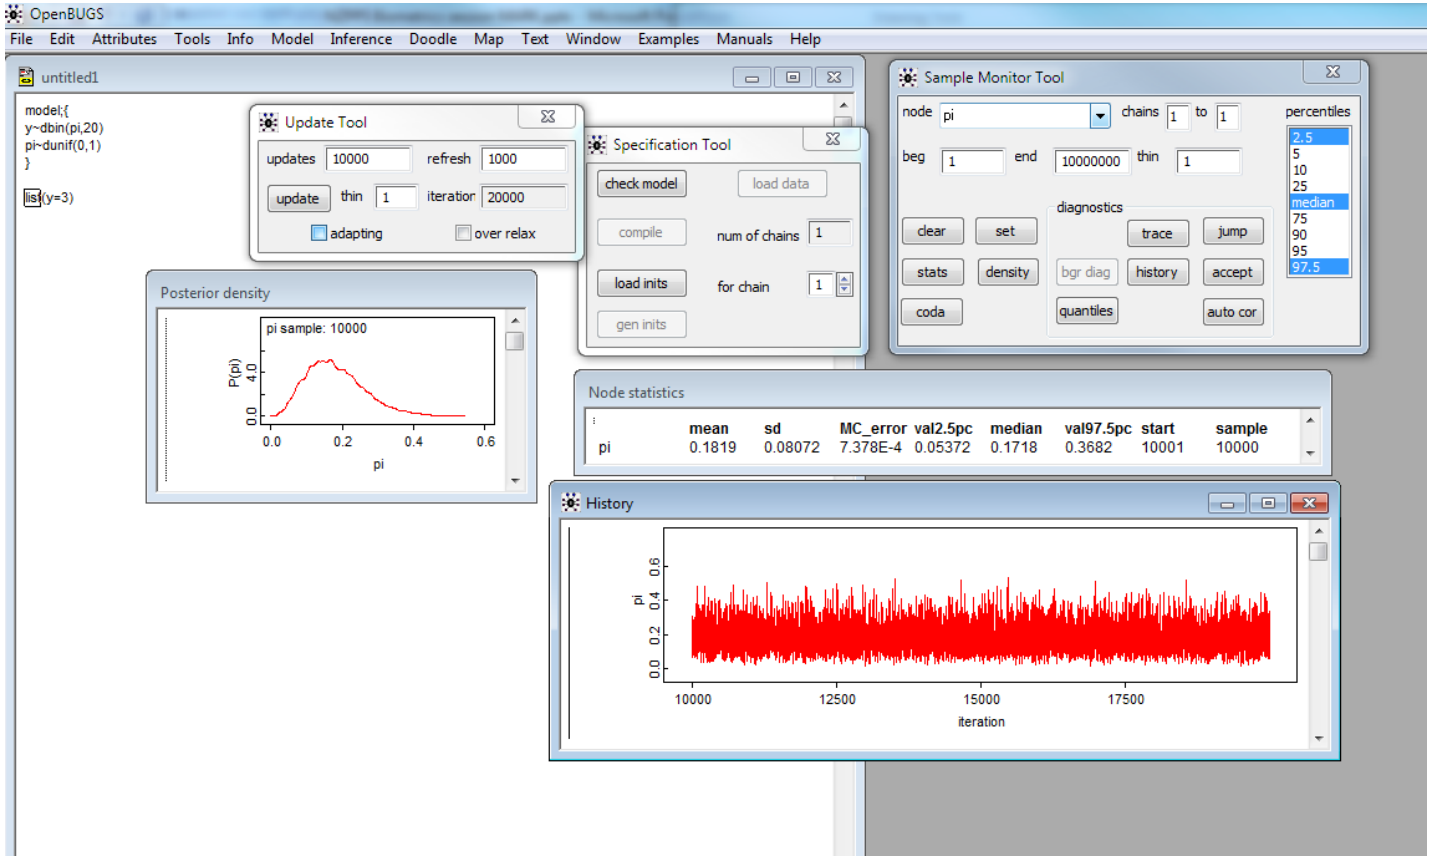

Bayesian Statistics

Criticisms

- Using prior knowledge introduces a subjective element to the analysis. Flat priors and sensitivity analyses are some methods used to address these concerns.
- With large samples the data will dominate and the Bayesian results should by similar to the frequentist equivalent.
- Can be overly complex and difficult to get the model to converge. More advanced algorithms have helped in this area.

Advantages

- Can produce more intuitive results e.g. Credible intervals vs. Confidence intervals, posterior predictive distribution.
- Doesn't rely on large sample approximations
- Forces you to think about the model and the sources of variation

Bayesian Statistics

http://twiecki.github.io/blog/2014/01/02/visualizing-mcmc/

1

Step Method

1) Metropolis-Hastings Algorithm
2) Gibbs sampling
3) NUTS (Hamiltonian Monte Carlo)

2

3