On Doing Research

To do …

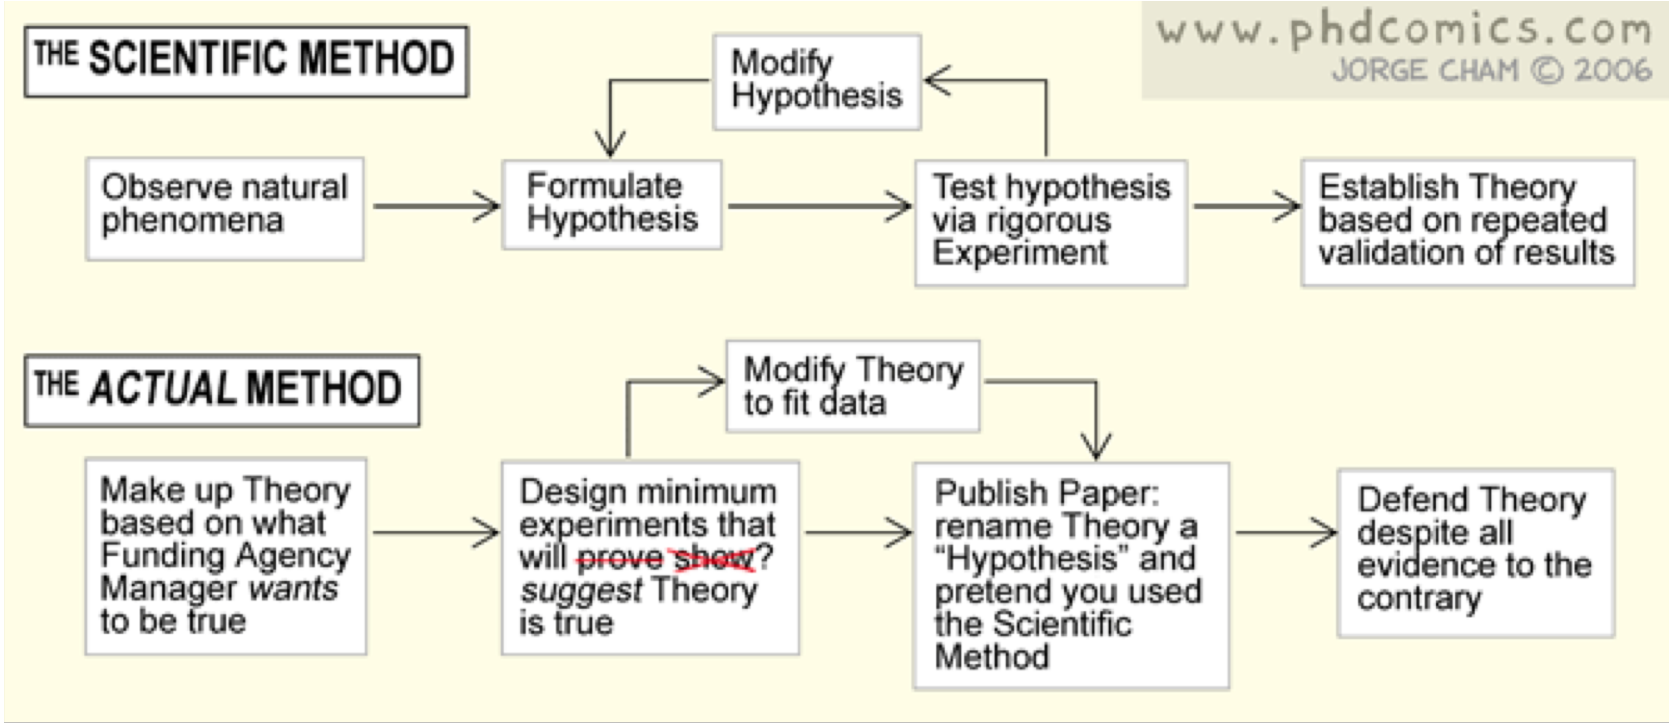

Doing research

- A paper is the result of long research process
  - Rarely final but a tiny window into an effort
  - The outcome of a cycle of activity from speculation to definition, experimentation, …
- Learning to do research, a piecemeal acquisition of a range of skills, learned by doing
  - The bigger picture of “the process of research” emerges from multiple, supervised research attempts
- Let’s look at the process of research, particularly the early stages

Doing research

- The beginning
- Shaping a research project
- Exploring related research
- Planning the effort
- Stating a hypothesis
- Defending a hypothesis
- Evidence
- All along, you and your advisor

The beginning

- Typically a moment of insight
  - *Why don't search engines do better spelling? Netflix in Cuba? C'mon …*
- When … many times when your brain is idling, or separate topics arise at the same time, or while talking to others, …
- First step is totally subjective
  - What do you choose to pursue among many ideas?
  - What if it doesn't pan out? It's ok, secondary directions are many times more interesting

Shaping a research project

- From topic to project depends on context
  - Experienced scientists aiming to write a paper on a subject tend to be very focused
- Two key questions to answer at the beginning
  - What is the broad problem to be investigated?
  - What are the specific initial activities to undertakes and outcomes to pursue?
- Short term goals give shape to the effort and helps training on elements of research
  - Planning, reading, programming, testing, critical thinking, analysis, writing and presentation

Shaping a research project

- When developing a question into a research project, what makes the question interesting?
  - Successful research is usually driving by a strong motivating example
- Sometimes have to decide to explore questions where work can be done rather than where we would like to work
  - Soccer playing rather than planetary exploration
- Risky choosing *a topic and advisor* focused on "is the most interesting topic on offer?"
  - To the exclusion of other questions equally important

Shaping a research project – to consider

- Is this the right advisor for you?
  - That's going to be a long, intense relationship, more soon
- A 'fashionable' topic – at most a minor factor
  - By the time you graduate could be passé
- Is the project the right kind of technical level?
  - Are you a hacker, do you have strong mathematical chops, …, does it fit your character (broad impact? Too speculative?)
- Project scope
  - Major breakthroughs are rare and risky; most research is incremental

Advisors are key to project scoping

- Stand sufficiently alone from other current work
- Yet still be relevant to the group's wider activities
- Open enough to allow innovation and freedom, but still with good likelihood of success
- Close enough to the faculty core area of expertise so she/he can tell about novelty, related work, etc.

Finding research literature/reading

- Finding relevant work
  - Visit websites of groups in the space; gives you an idea of conferences, co-authors, papers..
  - Follow up references in recent papers
  - Browse recent issues of conferences/journals
  - Use obvious search terms in Google Scholar
  - Discuss your work with others, similar problems often appear in other areas that you won't be aware of
- Reading
  - *We have covered this*
  - Be questioning and skeptical, yet not dismissive

Research planning

- In undergraduate, activities are determined by a succession of deadlines that give structure
  - Research has just one – completion
- So, scope the project and set deadlines
  - From a paper deadline, work backwards to figure out when you want to have certain pieces finished by
    - Helps to prevent the project from going unbounded
  - Figure out dependencies!
  - Then forward, time sequence of timelines for tasks

Hypothesis

- First steps
  - Identifying interesting topics, focusing on particular issues to investigate
  - A typical way, develop specific question you are trying to answer
  - The question requires an understanding, an informal model perhaps
  - This sets the framework for making an observation about the object being study – a hypothesis
- Key component of a strong paper – a precise, interesting hypothesis

Good hypotheses

- Hypothesis should be specified clearly and precisely and should be unambiguous
- May be important to state what is not being proposed
  - The limits of the conclusions
- Example from Zobel's
  - *p*-lists are well-known data structure used for a range of apps, as an in-memory search structure that's fast and compact
  - You develop a new structure, *q*-lists, asymptotically similar but you think superior in practice

Example …

- Hypothesis v1 – *q-lists are superior to p-lists* (x)
  - To be true in all apps, all conditions, all the time!
- Hypothesis v2 – *As an in-mem search structure for large data sets, q-lists are faster, more compact than p-lists* (!)
- Maybe a further qualification – *we assume there is a skew access pattern*
- Imposes a scope on the claim, others can find other apps that won't do or explore the behavior of q-lists under different conditions
- The hypothesis must be testable, it should be falsifiable
  - *Q-lists performance is comparable to p-list performance* (x)
  - Our proposed query language is relatively easy to learn

Defending a hypothesis

- Next for a strong paper – testing of the hypothesis, presentation of supporting evidence
- For presentation, construct an argument showing that evidence supports the hypothesis
- To construct the argument, imaging defending your hypothesis to a colleague that raises objects you have to defend against
  - If you can rebut objections, admit them; if you reasoned them away, include the reasoning
  - Basically, anticipate the reader's own objections

Evidence

- Broadly speaking, four types
- Analysis or proof – a formal argument that the hypothesis is correct
  - A common mistake, not all hypotheses are amenable to formal analysis (real world – people, systems, …)
- Model – a mathematical description of a hypothesis
  - There is usually a demonstration that it “fits”
  - In choosing a model, consider how realistic it will be, how many simplifying assumptions are being made

Evidence

- Simulation – an (maybe partial) implementation of a simplified form of the hypothesis
  - Wide range, from skeletal to detailed with artificial data
- Experiment - a full test of the hypothesis, based on an implementation of the proposal and on real – or realistic – data
  - Ideally done in light of predictions made by a model
  - Should be severe, looking for tests that will fail if the hypothesis is false

Evidence

- Different forms of evidence can be used to confirm one another
  - E.g., Simulation to confirm a proof's correctness
  - But not confused with one another
    - Running a program that implements an algorithm is not an experiment
- When choosing
  - Consider what you would need to convince your reader
  - Your community, at this time

To close - you and your our advisor

- Advisors are powerful figures in students' lives
- Among the closest of all your interpersonal relationships
  - Codified as "conflict for life"
- Look for compatibility in
  - Ideas: ambition level, vagueness level, goals
  - Management style: independence, hands-on vs. hands-off, structured vs. unstructured
  - Personality: humor, life perspective, etc

Ideally your advisor

- Feeds you with funding
- Feeds you with good problems to work on
- Guides you along the way to a good solution
- Teaches you all the unwritten skills of research, explicitly or implicitly, including writing, speaking, reviewing, grant-writing, etc
- Promotes you, internally and externally, for fellowships, jobs, committees, etc

Your part of the deal

- Your advisor is
  - Overloaded
    - Take notes, be frank
  - Ultimately an intellectual, and excited by ideas
    - Don’t wait to be fed, pick topics he/she cares about
- Your advisor is happy if
  - You save him/her time
  - You don’t create last-minute emergencies
  - You understand the high-level goals, and come up with things he/she didn’t think of
  - You learn on your own, and teach him/her
  - You don’t give up instantly

A Research Checklist

A research checklist

- Are the ideas clear and consistent?
- Is the problem worth the investigation?
- Does the project have appropriate scope?
- What are the specific research questions?
- Is there a hypothesis?
- What would disprove the hypothesis? Does it have any improvable consequences?
- Are the premises sensible?
- Has the work been critically questioned? Have you satisfied yourself that is sound science?
- How are the outcomes to be evaluated? Why are the chosen methods of evaluation appropriate and reasonable?

A research checklist

- Are the roles of the participants clear? What are your responsibilities? What activities will others undertake?
- What are the likely weaknesses of your solution?
- Is there a written research plan?
- What forms of evidence are to be used?
- Have milestones, timelines and deadlines been identified?
- Do the deadlines leave enough tie for your advisor to provide feedback on your drafts, or for colleagues to contribute?
- Has the literature been explored in appropriate depth? Once the work is mostly done, does it need to be explored again?