THE UNIVERSITY of EDINBURGH

HERIOT WATT UNIVERSITY

Robotics Research Proposal (2)

J. Michael Herrmann

University of Edinburgh

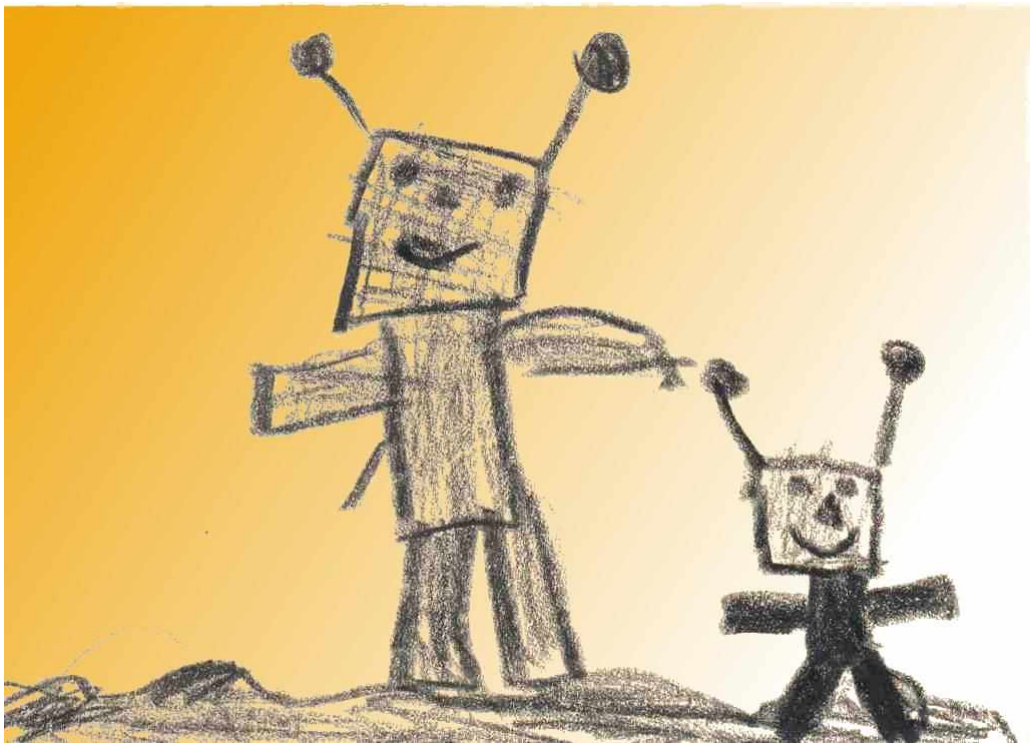

RRP: What to do

- Establish aims and objectives of the project
- Establish hypothesis and evaluation
- Break project into work-packages
- Submit full proposal (deadline: 4pm 12/4/2019)
- Same submission procedure as for RRR
- Use turnitinuk.com

RRP: Guidelines for Writing a Research Proposal

- A good proposal will provide a convincing case for the high quality of the proposed research.
- It will show an awareness of relevant prior work and include a clear statement of the problems and hypotheses to be addressed and why they are important.
- It must also make clear exactly how the methods used to research those hypotheses will yield interesting results. There are many ways in which one might structure the material.

DRPS: RRP (IRP)

Assessment
The assessment will come from one piece of submitted work: a full research proposal, including background, motivation, and a description of the research methodology and expected outcomes. A good proposal might be organised as follows:

- Purpose: a statement of the problem to be addressed.
- Background: a short description of how previous work addresses (or fails to address) this problem.
- Methods: A description of the methods and techniques to be used to test the hypotheses, indicating that alternatives have been considered and ruled out on sound scientific grounds.
- Evaluation: Details of the metrics by which the outcomes will be evaluated.
- Workplan: A timetable detailing what will be done to complete the proposed project, and when these tasks will be completed.

As a guide, a good proposal might be organised as follows:

Purpose: a statement of the problem to be addressed. This should include arguments as to why solving the problem is important; e.g., because it will enable certain applications, or lead to interesting scientific discoveries.

Background: a short description of how previous work addresses (or fails to address) this problem, leading to a rationale for the hypotheses that you intend to test, and a convincing argument about how that hypotheses might solve the problem.

Methods: A description of the methods and techniques you intend to use to test your hypotheses (e.g., data analysis procedures, experimental design etc), indicating that alternatives have been considered and ruled out on sound scientific grounds.

Evaluation: Details of the metrics by which you will evaluate the outcomes of your research; e.g., by comparing the output of your system with some gold standard, or with the ways in which humans perform a task, etc.

Outputs: A description of what the outputs of the projects will be: e.g., these might include an extension or change to some existing theory or to some piece of software, some new data (e.g., annotated linguistic data), and so on.

Workplan: A timetable or research plan, detailing what will be done to complete the proposed project, and when these tasks will be completed by.

The proposal may also include material that would count as the Introduction to the MSc thesis itself, and/or the literature review.

Hypotheses

- A potential explanation for a phenomenon
  - **A** (facts) can be explained by **B** (reasons)
  - If **B** then **A** (tentatively)
- Hypotheses vs. theories
  - A hypothesis can be stated without evidence
  - Evidence turns a hypothesis into a theory
- What is a good hypothesis?
- How do arrive at a hypothesis?
- Is it actually necessary to formulate a hypothesis?

Hypotheses in Informatics / Robotics

- *Hypotheses/claims* often not stated
  - *except* in *theoretical* work
  - leads to confusion and misunderstanding
- If *claim not clear* then this should be *criticised*
  - same if claim is *strong* and is *not proven*
- *Evidence* may be *theoretical* or *experimental*
- *Objective* may be to *identify a hypothesis* for subsequent evaluation

Hypotheses

- Philosophy is not needed specifically in robotics, but a bit of it is always needed
- “No amount of experimentation can ever prove me right; a single experiment can prove me wrong.” (attributed to Einstein)

What is a good hypothesis?

- Testability, falsifiability
- Parsimony, Occam’s razor
- Scope, applicability of the theory
- Fruitfulness, likelihood of the outcome, achievable information gain, impact
- Conservatism, probability of acceptance

Falsifiability

- K. Popper claims that statements that are not falsifiable are not scientific
- D. Hume: Positive evidence adds merely likelihood, but does not prove anything
- A scientific statement is neither true nor false, but is acceptable as long as has not been falsified
- Examples (to be discussed)
  - “All swans are white.” (refuted by Willem de Vlamingh, 1697)
  - “There is no bound to the number of angels that fit on the tip of a needle.”
  - “It is impossible to perceive a (homogeneous) green Red.” (I. Kant)
  - “It is possible to construct a robot that does X.”
  - “Spontaneous generation does usually not occur.” (L. Pasteur)
  - “If c is the length of the longest side of a flat rectangular triangle and a and b are the lengths of the other two sides, then $a^2+b^2=c^2$.”

Falsifiability: Criticism

- “Exceptions prove the rule”
- Explanatory power: A statement with more “exceptions” can be preferable to a competing statement it explains also more data
- Statistical statements
  - “Controller A is better than controller B”
- Testability vs. Falsifiability
- Scientific programmes and successes
- Paradigm shifts (Th. Kuhn)

How to arrive at a good hypothesis?

- Specificity: What is goal of your research?
  - Why this goal has not yet been reached?
  - What do you believe will help to reach this goal?
  - Is this the best way to reach this goal?
- Side-effects: Scope, impact
- Formulation may still be difficult.

How to arrive at a good hypothesis?

Examples (not necessarily good ones)

- “Birds form flocks in order to save energy.”
- “Birds form flocks in order to improve navigation.”
- “Robot swarms are useful in disaster areas if they have the ability to self-localise.”
- “The swarm acquires the ability to self-localise via self-organisation”
- “Robots can cluster into a single cluster without central control or central sensing.”

How to arrive at a good hypothesis?

Variants:

- "Goal can be reached without condition **C**"

  Improve generality or increase scope
- "No clear goal, it's more like a dream"

  Try to find out why it doesn't work
- "I want to understand robot swarms better"
  - What is it you don't understand?
  - Try one of the above
  - Try to be more critical, disprove someone else (they may not understand it well either)

“Data-driven research”

- If there is the perception, that currently used hypotheses are not representing important practical questions, then research may try to extract hypotheses from data
- Helps to escape “baroque” theory frameworks
- Information ≠ **Data** ≠ Knowledge
- Short-cuts from data to decisions
- More general hypotheses will still be needed

Towards a workplan

- Do the discussed approaches show any gaps, or unused potential?
- Are combinations with other approaches promising?
- Why has the potential not been used before?
  - it just became visible
  - they suffered from limited resources
- You will need (to be able) to reproduce existing work. This is a first step.
- now: Hypothesis, Realisation, Evaluation (as discussed before)

RRP: legal, social, ethical and professional issues

Questions: Are you happy with all aspects of the work?
If you are happy, will everybody* else be happy too?

*this may exclude competitors

Are there aspects that deserve particular attention?

A number of legal, social, ethical & professional issues can occur: E.g.

- Privacy issues (Databases, learning from data, knowledge management)
- Ethics of human and animal experiments (HCI, HRI, Neuro- and cognitive science)
- Weaponry (Intelligent Robotics)
- Non-disclosure (Industry collaboration)
- Legal issues (Natural Language Processing, image processing)
- Best interests of client and providers (Computer Systems)
- Consistency with the public interest
- Integrity and independence in the professional judgement
- Fairness and support to colleagues

RRP as part of a thesis project

- Literature review
- Specification of a direction, goals and methods
- Justification of the approach
  - filling a gap that was identified in the literature
  - similar to examples from the literature
  - a new combination of existing approaches
  - application of a existing approach to a new domain
  - extension, generalisation, removal of assumption
  - improvement of existing approaches

Some considerations

- With the submission of the RRP you have finished work worth 40% of the project
- The more results have obtained before submitting the Proposal, the more realistic your plan will appear
- The more work you have done by the end of this term the better can your project expected to become.
- The content of the proposal is basically the same content as the thesis
- Except that results will be just plans for the moment) (Alternative plans in the proposal will be points in the discussion of the thesis)

Robotics Research Proposal (RRP)

- Submission last week of term (4pm 12/4/2019)
- Marked by the same criteria and procedure as RRR
- Mark will count towards project mark (30%)

- Pace yourself
  - Leave time for feedback and correction
  - Self-assessment against marking criteria

- Meet with your supervisor regularly
    - If they are unavailable, keep contacting them
    - Talk to Ras and PhD students in the group
    - If problem persists, contact me:
      michael.herrmann@ed.ac.uk

Introduction (M. v. Rossum)

- What problem are you working on?
- Why is this an important / interesting problem?
- What is the core idea of your solution?
- Which questions/hypotheses are you trying to answer with your work?
- What is novel/original about your solution?
- How are you going to test if it works?
- What are the main contributions / salient points in your project?
- Brief overview the rest of the proposal

- Be very clear about which ideas are your own and which are not, especially if you are working in a groups; cite all sources

Background (Research review)

Which fields of research are closely related to your work? (should be 2-3)
What are the most important (highly cited?) publications in those fields?
How are they (are going to be) related to each other?

How is your work similar to those publications (update RRR)?
– are you borrowing ideas / motivation / algorithms?
– are you using similar datasets?
– similar evaluation framework?
Aspects that were less relevant in RRR (partly in **Methods**)
– what differentiates your work from the prior work?
– what motivated the difference in approach?
– what aspect of the problem are you hoping to do better?
– will you be comparing your performance to the prior work?
– summarise related work using a common vocabulary

Methodology: provide a high-level outline of your solution
– what tools are used in the major steps of your project: e.g. pre-process, extract features, measure similarities, ...
– describe each phase in a separate sub-section:
– think in terms of what goes in and what comes out, but it is not necessarily the best way to describe it that way
– be clear about the purpose of each step and include a justification
– discuss design decisions and explain why you chose A over B
– describe possible variations on the approach / parameters that will affect performance
– use equations and diagrams to specify and illustrate your ideas & assumptions
– use standard terminology and be consistent; don't use synonyms
– define all symbols and use them consistently in equations / diagrams / text
– include pseudo-code for complex algorithms, but keep it brief and high-level
– usually you would not include class diagrams, robot construction plans etc. (put them into appendix and point to them in the text)
– devote a special sub-section to summarising the steps / phases

Work plan

Describe the datasets you will be using or acquiring
– what experiment did/does the dataset come from?
– what are the summary statistics (dimensions, number of instances)
– what sort of pre-processing was/will be done
Describe evaluation methodology (perhaps already in **Methods**)
– what metrics will you be using and why?
– what does the ground truth look like and how was it generated?
For each hypothesis:
– describe the design of your system (or difference to previous cases)
– describe the baselines you will be comparing to
– state the hypothesis precisely: is your system expected to be faster? more accurate? under what conditions?
– describe the main experiments you are going to carry out
– how many experiments are needed for statistical significance
– discuss whether the system is ready to use or whether parameters need to be adapted
– describe the expected outcome for each hypothesis
– Include a special sub-section to summarise all hypotheses

Discussion

Revisit the claims/statements you made so far.

– Does your workplan meet expectations about the outcome that you have formulated in earlier chapters?

– Are any major surprises expected? Will all the resources, tools, software etc. be available? Discuss fallback options

– What peculiarities did you encounter in working with the robots, algorithms etc.? Can these problems be discussed away or will the require further work?

– What are the limitations of your approach? Where would it fail?

– Critically compare your approach to prior work: Discuss range of applicability of your system and of the baseline (i.e. your solution may be better only in certain cases)

Conclusions

Re-cap the hypotheses you tested and the main results (chapter 4)

– What are the major lessons learned? What should the reader take away from this thesis?

– What would be different in your approach if you were to do the project again?

– Future work: imagine you had a year to continue working on this project

– which questions would you focus on?

– what approaches would you consider?

– what resources would you need?

– having done the project, what do you see as the biggest challenges in the field?

Back matter

Bibliography: 20 - 50 citations is a reasonable number

Appendix: Bits of code, class diagrams, directory structures, study questionnaires, long tables and tables of graphs

Argument Clinic

Why arguing?

- Facts often speak for themselves, but complex scientific facts usually don't.
- Ideally, the results follow from the assumptions and data, by predicated logic, but practically you cannot express all details in this way.
- A good argument avoids logical fallacy.
- What else qualifies an argument as “interesting”?

  Generality, relevance, elegance, topicality, timeliness, attractiveness, understandability

see also http://staffhome.ecm.uwa.edu.au/~00043886/humour/invalid.proofs.html

Fallacies and why they are here to stay

- Prove a general statement by example: An example does often establish meaning or relevance.
- Referring causes/reasons that are not clear: To an expert the reasons may actually be clear.
- Arguing without arriving at a relevant conclusion: You may need to circumnavigate a cliff not visible to the landlubber.
- Association: Relating a claim or approach to a fashionable buzzword: What else in the era of search engines?
- Proof by authority/funding/application: Impact and success may be preferable to pure understanding.

Strategies (Mix to taste!)

- Limit yourself to a small range, argue overly clearly about justification and correctness.
- Report on what you did, avoiding general implications.
- Combine two lines of research or two ideas, and show that this has advantages and does not cause a problem for either.
- Give a broad background, make clear where your work is important because of the background (although it is just a minor step).
- Do the same as everyone else and show you are doing it better (or at least slightly differently).
- Present a reasonably good idea as something novel (new beyond comparison).
- Try to derive the most general statement that is defensible.
- Get closer to an application that makes a difference, increase TRL

Characterisation of research

- Falsifiability → Invulnerability
- Reproducibility → Veracity
- Creativity → Elegance
- Accessibility → Topicality
- Quantification → Comparability
- Qualification → Novelty
- Essentiality → Generality
- Relevance → Applicability

How much of this is expected in a Master's thesis?

- There are exceptional criteria in marking
- An opportunity to try out one or the other strategy
- Continuously improve your style
- Thesis will be the main result of your research: Composition rather than merely writing-up
- Be proud of yourself: If your results appear trivial to you, then you really did understand something and you are ready to move on.

Sequentiality

- Sometimes the argument is rather a (directed) network than an ideal linear chain of consequences
- To express this an an understandable text suppress some of the secondary thoughts, otherwise use
  - a detailed document structure several levels of sections
  - cross-references and references to literature
  - footnotes
  - unambiguous grammatical reference
  - diagrams and schemata
  - a mild amount of repetition
- Have your text read by your fellow students (there is no relative marking of theses)

Conclusion

- RRP should not be an obstacle to get started with the work on the thesis
- Make the project yours!