Goldwater Research Essay Q&As

Providing Some Guidance on Writing the Research Essay

Kerry Karukstis, karukstis@g.hmc.edu
Bert Holmes, bholmes@unca.edu
Paul Zorn, zorn@stolaf.edu

These slides will be posted on the Goldwater web site under the "Open House Presentations" tab in the lower left corner of the opening page.

In this Open House, we will:

- Review the reasons for requiring a research essay in the application
- Provide some guidance on selecting the research to be described in the essay
- Address some questions on the format of the essay
- Provide examples of how to address the goals of the research essay
- Review an example of a research essay for its strengths
- Examine an AI meta-analysis of research essays from several successful applications that generated a guide to developing a research essay

Why does the Goldwater application include a research essay?

Reviewers are looking for **evidence of the student's motivation, background, experience, skills, and interest in pursuing a research career**, as demonstrated in the description of a completed or proposed research project. Reviewers also looks for typical characteristics of successful researchers such as determination, creativity, resilience, etc.

Strong essays describe:

- The importance of the research project to the discipline.
- The student's specific role and intellectual contribution to the work.
- The skills the student brought to the research from prior research or course experiences and applied to the current research.
- The new skills gained from participating in the project.
- Examples of how the student "thinks like a scientist". These might include problem solving during experimentation or computation, data analysis, or interpretation of project results.

https://goldwaterscholarship.gov/developing-the-research-essay/

Let’s start with some questions on selecting the research project to describe in the essay.

Q: If an applicant has conducted more than one research project, should the research essay summarize all of the projects?

A: Focus on one project and focus on the project in which the student (1) demonstrated the greatest **ownership**, (2) achieved the most **significant results**, and (3) clearly showcased his/her **ability to think like a scientist**. The project will likely be the one in which the student spent the most time (but not necessarily). Other projects can be briefly summarized if the students thinks that would be informative.

Be sure that one of the letters of recommendation is from the research mentor for this project.

Q: If an applicant has conducted more than one research project, should the research essay focus on the most recent project?

A: The project does **not** need to be the student's most recent. Use the same criteria noted on the previous slide to select the project to address. In other words, chose the project in which the student (1) demonstrated the greatest **ownership**, (2) achieved the most **significant results**, and (3) clearly showcased his/her **ability to think like a scientist**.

Q: Can the research essay focus on a course project?

A: Yes, but this option is probably best if the applicant has no independent research projects outside of courses. The course project should be ***discovery-based*** (i.e., the results are not known to the student in advance) and ***require substantial individual student participation***.

Again, be sure that one of the letters of recommendation is from the course instructor who assigned the research project.

Q: When should a student submit a research proposal rather than a research essay?

A: This option is probably best if the applicant has no independent research projects and no research projects in courses. Be sure to address all the elements of strong research essays as outlined on slide #3.

It would be ideal if one of the letters of recommendation is from a faculty member who can speak to the ***originality*** *of the proposal, the* ***scientific merit*** *of the question being addressed, and the* ***capability*** *of the student to conduct the proposed research.*

Let’s address some questions on formatting and developing the research essay.

Q: Is it OK to use equations in the essay?

A: **Yes.** Equations are appropriate, even necessary, when they communicate important relationships in a compact manner. Essays on mathematical topics, for instance, will usually include pertinent equations. Essays in theoretical physics and computational chemistry are other areas where equations may be essential to convey particular relationships. Be sure to identify all variables and explain the context carefully.

Putting one’s work in context is important in all disciplines, but perhaps even more so for highly sophisticated mathematical research. As one reviewer shared, a useful description of the research that makes it accessible to the reader is absolutely necessary to accurately assess the value of the work.

Q: Is it OK to use figures in the essay?

A: **Yes.** As appropriate, utilize graphs, tables, and figures in the essay to explain or clarify results. Keep the number of such items to a minimum. Use captions to guide the reader in interpreting graphs, tables, and figures.

Q: Should the applicant use headings to demarcate the various sections of the research essay? If so, are there preferred headings to use?

A: Using headings is an ideal way to organize the essay and help the reviewer to follow the student's train of thought. There are no set preferred headings. Use the terms that provide structure and clarity for the student's writing. Only a few headings should be necessary in a short essay.

Q: Can you suggest some preferred ways for the student to introduce and conclude the research essay?

A: A short opening paragraph can present the broader picture to the reviewer, specifically the *importance of the research and the motivation for doing the work*.

A final paragraph might discuss either (1) *next steps* for the project if the work is not yet complete or (2) *possible new projects* that the current research has inspired.

Q: Can you give some examples of how a student can describe their own contributions to a research project?

A: Use the word "I" and provide examples of activities like:

- Taking initiative and working independently
- Proposing solutions and strategies to roadblocks or unexpected results
- Actively contributing ideas
- Making decisions and being responsible for the outcome
- Remaining committed to the project despite setbacks

Reviewers understand that research is collaborative and that students will need to gain expertise in new areas and techniques. Avoid language such as "Working with a graduate student, I…" and "My lab partner and I…" and use phrases such as "I became proficient in …." and "I contributed …"

Q: The Goldwater instructions on the research essay mention discussing "the skills you brought to the work from prior research or course experiences" and "the skills you have or expect to obtain from participation in the project." Can you give some examples of what is meant by "skills"?

A: The skills that reviewers are looking for might be categorized as follows:

Technical skills - scientific techniques associated with conducting the research involving instrumentation, software, writing, presentation

Research skills - critical thinking, hypothesis formulation, data collection and analysis, evaluation of sources, communication, time management, attention to detail, problem-solving, project management, collaboration, leadership of a team, adaptability, and the ability to plan and organize research effectively

Q: How can a CR assist applicants in developing their research essays?

A: Feel free to share the advice presented here with applicants. Also encourage the applicant to:

- Start drafting the essay well before the deadline and write multiple drafts.
- Share drafts with the research mentor.
- Seek additional feedback from others such as a STEM faculty recommender, a STEM writing consultant in the campus' Writing Center, the Campus Representative, the Goldwater selection committee, etc.
- Pay close attention to spelling and grammar.
- Ensure that the final essay maintains one's individual voice.

A Campus Representative recently posted the following research essay on the National Association of Fellowships Advisors (NAFA) listserv.

While the student author of this essay was selected for a Goldwater Scholarship, keep in mind that this essay is one example of how to organize the essay and demonstrate the student's contributions to the project.

As we have noted numerous times at our Goldwater Open Houses, reviewers evaluate applications ***holistically****, giving balanced consideration to all elements of the application. Applicants should write their research essay to reflect their own voice.*

Simple heading

Opening paragraph to describe:

- the importance of the research
- the motivation for doing the work

Learning Deep Goal-Parameterized Robotic Skills from Demonstration

Introduction

It is often difficult to manually program even simple skills for robots: commanding a robot requires sending a certain precise current to each motor, which is especially complicated when a set of motors within an arm must work in unison to achieve a complicated movement. Learning from Demonstration (LfD) is a promising direction of research for robots to learn real-world skills directly from observing demonstrations. LfD allows non-expert operators to program skills simply by demonstrating them many times. Furthermore, these learned skills are more *general*: they are able to handle slight variations of a task; such as if an object to be picked is slightly misplaced.

However, many contemporary LfD approaches train skills that are unable to target a specific goal from many possible choices. For instance, such approaches are unable to target a specific button within a grid. Instead, these approaches would train a different skill for every button. This requires a lot of data (approximately 30 minutes of demonstration data per skill [Zhang et al., 2017]) that is usually infeasible for real-world situations.

* Multiple uses of the word "helped" suggest a collaborative endeavor. If this is not the case, the student should reword. The student might consider some stronger language, such as "along with others I proposed a method…".

To combat this issue, I helped propose a method that learns skills that are parameterized by a goal-parameter (τ) such that altering τ correctly alters the skill. In the button-pressing scenario, instead of training a new skill for each button, we train *one* general skill that adapts itself depending on where the button is (τ), much like how a human might learn this skill. This enables a robot to press a new button (that it hasn't seen in demonstrations) simply by being given the button's goal-parameter. *

This work directly builds on recent successful approaches to LfD. Zhang et al. [2017] and Levine et al. [2016] used Deep Neural Networks (DNN's) to approximate functions mapping from a robot's sensor input to the expert's action. These DNN's are able to learn complicated real-world tasks (such as block-stacking) from demonstrations. I proposed and helped implement a novel DNN architecture that builds directly on this success by adding a goal-parameterization (τ) variable as an input to Zhang et al. [2017]'s DNN architecture. I then designed and helped run experiments on a variety of simple tasks and representations for τ to evaluate our method's performance empirically. * *

Description of how the research contributes to and advances the discipline

Example of student's creativity and contribution to the project

* Multiple uses of the word "helped" continue to suggest a collaborative endeavor. Research mentor should clarify the student's exact contributions.

Another helpful heading

Actively designed experiments

Figures give the reader a concrete illustration of the three experimental set-ups

Methodology

In an effort to evaluate our algorithm, I designed several experiments in three separate domains: a 2D simulation of the button-pushing task, a 3D button-pushing task with a physical robotic arm and finally a 3D peg-insertion task with a different robot arm.

Figure 1: Views from our experimental setups

2D Button Simulation

Our experiments in this domain aimed to answer the following questions:

1. How does our method's performance compare to the state-of-the-art?
2. How does the representation chosen for the goal-parameterization (τ) affect our method's performance?
3. How does our method's performance change with the number of different goal-parameters (values for τ) provided during training?

To answer these questions, I designed a 3$\times$3 grid of blue squares representing buttons and a black circle representing the agent, shown in Figure 1. I then collected 100 expert trajectories for each square where the agent began at a random position along the right wall and followed a randomized path to the specific square. I helped train our DNN on random subsets of squares in our grid. For every random subset, we evaluated our DNN's median performance on 100 trials for each button. A trial was counted as a success only if the agent stops at the correct blue square. *

I helped perform the above experiment for various versions of our DNN. One version did not take τ as an input parameter and was thus equivalent to the state-of-the-art architecture from Zhang et al. [2017]. Another version used an unstructured representation for the goal-parameter to study how structure in the choice of representation for τ affects our performance. Specifically, it used a one-hot vector where a goal corresponded to a randomly-chosen index of a nine-dimensional vector (as there are nine possible goals). Finally, I tested version of our model that used the structured representations for τ, specifically the button's pixel location within the image as τ or its row-column index pair [for example $(1, 1), (1, 2)$, etc.]. *

Conducted tests of proposed model

This section details numerous ways that the student contributed to conducting the investigations.

3D Robot experiments

In this domain, I aimed to investigate whether our method would work robustly on real-world robotic tasks. I used a KUKA robot arm to press buttons on a 3D, 4×4 button panel pictured in Figure 1. Similar to an experiment from our 2D Button Simulation Section, I parameterized our button-grid with a row/column tuple of the button's location on the grid. For training, I collected 100 trials of the robot's end-effector beginning at a random position and following a straight line to the specified button. The end-effector's final position was varied with noise distributed normally such that the robot would press the button differently each time.

For this experiment, I used specific subsets of buttons that had been found to generalize well in our two-dimensional simulation. After having trained our DNN on the data, I evaluated this learned policy by averaging three attempts of the robot attempting to press the button.

I helped repeat the same experiment on a peg-insertion task (depicted in Figure 1) to study whether our method can perform a task that requires significantly more precision.

Results

2D Button Simulation

Figure 2: Results from the 2D simulation domain

From Figure 2 above, one can infer the following answers to questions posed in the 2D Button Simulation Methodology Section:

1. Our method performs significantly better than the existing state-of-the-art. The 'no tau' curve's success percentage drops to a relatively consistent 0 for any number of goals greater than 1. This is almost certainly because the state-of-the-art method doesn't take τ as an input and thus cannot target specific goals when trained on more than one goal. Instead, it learns to average between the goals seen during training.
2. The representation chosen for τ affects our DNN's performance rather significantly. Both the row/column index parameterization (index tau) and pixel parameterization (pixel tau) achieved much higher success percentages when trained on fewer goals than the one-hot vector parameterization (onehot). This is probably because τ varies inconsistently for the one-hot parameterization and thus the DNN is only able to target goals already seen during training (hence, the observed straight-line trend).
3. Regardless of the structured parameterization chosen, our DNN's performance improved with the number of goals seen during training. For the index and pixel parameterizations, the average success percentage reached 100 after training on just 4 of the 9 possible goals in the grid.

Concluding paragraph detailing next steps to continue the project.

3D Robot experiments

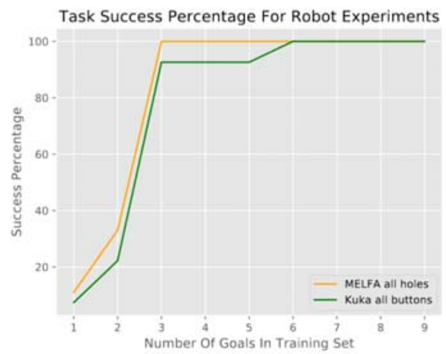

Figure 3: Results from our robot experiments

Example of student interpreting research findings.

Figure 3 illustrates that our DNN is able to solve both targeted button-pressing and peg-insertion tasks remarkably well on robots. For both tasks, our DNN is able to learn to generalize to target all possible goals without the need to train on all 9 goals. Interestingly, our method performed better on the peg-insertion task even though it required more precision and had a smaller amount of training data. I hypothesize that this is because there was more noise in the robot's motion for the button-pressing task, leading to a more imprecise policy that would narrowly miss specific buttons. Indeed, we observed this qualitatively.

Future Work

I hope to deepen this line of work in the future. Specifically, I hope to extend our idea of goal parameterization to other LfD frameworks; such as that adopted by [Ding et al., 2019], which has shown encouraging results in simulation. I hope our extensions to such methods will enable us to represent more complex, desirable parameterized skills (such as throwing a basketball into a hoop at a specific location) than button-pressing or peg-insertion.

Key references

References

Yiming Ding, Carlos Florensa, Mariano Phielipp, and Pieter Abbeel. Goal-conditioned imitation learning. *CoRR*, abs/1906.05838, 2019. URL `http://arxiv.org/abs/1906.05838`.

Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.*, 17(1):1334–1373, January 2016.

Tianhao Zhang, Zoe McCarthy, Owen Jowl, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

A recent post on the National Association of Fellowships Advisors (NAFA) listserv provided a Generative AI meta-analysis of a number of examples of research essays from the institution's successful Goldwater applications. The following slides detail the resulting guidance for preparing the essay.

A Goldwater Scholarship research essay should showcase the applicant's understanding of their research experience, their ability to communicate complex scientific ideas clearly, and their intellectual curiosity. Based on examples and best practices, the following key points should be covered:

1. Clear Description of the Research Project Purpose and Significance: Begin with a concise explanation of the research question or hypothesis and why it is important. Highlight its relevance to the field and potential real-world applications.

Context: Provide background information to establish the scientific foundation and situate the research within broader scientific trends or challenges.
Specific Role: Clearly state your role in the research, distinguishing your contributions from those of collaborators or advisors.

2. Explanation of Methods and Approaches Methodological Clarity: Describe the methods and approaches used in the research, ensuring they are detailed enough to demonstrate your understanding but clear enough for a scientifically literate audience.

Tools and Techniques: Mention specific tools, technologies, or techniques you employed and how they were critical to the project.

3. Results and Outcomes

Summary of Findings: Clearly explain what the research uncovered or achieved, whether it produced conclusive results, was exploratory, or remains ongoing.
Analysis: Demonstrate an understanding of how the results were analyzed and their implications.

4. Impact and Broader Implications

Relevance: Discuss how the research contributes to the field and its potential implications for future research or applications.
Connection to Goals: Highlight how this research aligns with your academic and professional goals as a future researcher.

5. Intellectual Engagement

Challenges and Solutions: Reflect on the challenges encountered during the research and how you addressed them, demonstrating problem-solving skills.
Insight and Growth: Share what you learned through the research process and how it has shaped your perspective as a scientist.

6. Concise and Accessible Writing

Technical Accuracy: Use precise scientific terminology, but balance it with accessibility for readers who may not specialize in your field.
Logical Structure: Ensure the essay is well-organized, with a logical progression from research description to broader significance.
Engaging and Clear Language: Avoid jargon-heavy explanations. Strive for clarity, concision, and elegance in writing.

7. Future Directions

Next Steps: Briefly outline future research directions for the project or related questions that the research has inspired.
Integration with Career Goals: Tie the research to your aspirations, emphasizing how the scholarship will help you continue to contribute to the field.

Focusing on these elements ensures the essay effectively communicates both your technical expertise and your potential as a future leader in STEM fields.