Role-Playing Paper-Reading Seminars

Alec Jacobson and Colin Raffel
March 17th, 2021
colinraffel.com/blog

The paper-reading seminar is a common format for advanced topics courses in Computer Science graduate programs. Over the span of the course, students read a series of academic papers published on a particular topic (for example, Deep Learning with Limited Labeled Data). For cutting-edge topics close to the frontier of research, standard texts may not exist and academic papers are the primary available references. Some niche areas may not be mature enough to develop a lecture-based course curriculum around. Complementary to course topic-specific learning objectives, paper-reading seminars aim to achieve learning outcomes related to the soft skills of technical reading, public presentation, and critical discussion.

1 The Standard One-to-Many Format

In its near-universal incarnation, each class session consists of a single student (or subgroup of students) giving a presentation that summarizes a single paper to the rest of the students. This one-to-many format is a straightforward division of student work and requires little instructor work outside of class time. We believe this standard format has various drawbacks, which we describe below.

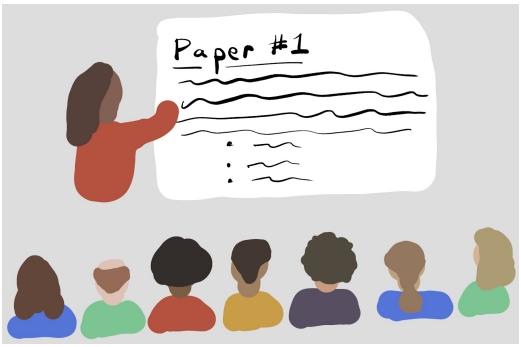

Students are only directly participating in the course in sessions when they are presenting, which may only happen once or twice each term. It is easy for students to disengage during the presentations of papers for which they are merely audience members. This results in their missing out on a large position of the course's content. To combat this, some instructors will *mandate* engagement via weekly quizzes or requiring each audience-student to ask a question of the presenter. While this can be successful in creating discussions in the first few sessions, students focused on maximizing marks while minimizing work may realize that vacuous questions such as “How efficient is this method?” can be asked near universally of any paper without having read the paper or listened to the presentation. Recognizing that audience

attention has diminished, as the term progresses presenter-students may themselves devolve into offering only shallow summaries of the selected paper void of critical analysis. Often the best-case scenario of the standard one-to-many paper presentation format is a tour of a list of papers understood at the most basic level.

The standard format also misses an opportunity to maximize supplemental learning objectives. Presenting a summary of an existing paper does provide students practice giving oral presentations, but this only happens once or twice during the term and evenso in a low stakes environment with little opportunity for feedback.

Certain skills are important for a research-focused graduate student to develop but are not directly taught through dedicated coursework in many Computer Science graduate programs. In particular: how to conduct peer-review of a scientific paper, how to generate new research ideas, how to rapidly reimplement a prototype of an existing paper, how to assess which aspects of an academic work are ready for practical use, and how to see beyond technical contributions of scientific work to its positive and negative societal impacts.

It is generally assumed that students will develop these skills on their own over the course of their research career. They are often expected to do so with limited formal instruction and feedback. Neither the classic lecture-assignment format nor the standard one-to-many paper reading format readily offers opportunity for learning these skills.

2 The Many-to-Many Role-Playing Format

Over the past few years, we — Colin Raffel at University of North Carolina and Alec Jacobson at University of Toronto — have taught paper-reading seminars that use an entirely different format centered around weekly rotating "role playing". This format has evolved from an initial design by Alec and Eitan Grinspun at Columbia University. In this format, many (or all) students in the class cooperatively present each paper. Students read the same paper, but each student takes one of seven specific roles for that week. The role defines the lens through which they read the paper and what they bring to the discussion that week. Students cycle through roles throughout the course.

For example, one role is the "Reviewer" who completes a full conference-style peer-review of the paper. The student playing this role justifies their decision to accept or reject the paper and

what changes they would require for acceptance. We have found that having students role-play in this way makes them engage much more deeply with the papers, maintain that engagement throughout the term, and provides explicit opportunities to develop important skills that they otherwise would have no formal instruction on.

As a positive side effect, the class is more rewarding for both the students and the instructors. In the standard format, the instructor may be reduced to a validator of the paper-summary presentation or as a depositor of knowledge through the instructor's post-presentation comments. The standard format fits squarely into the "banking model of education," where the audience-students are containers into which knowledge is deposited [freire1968]. In contrast, in the role-playing seminar the instructor brings their own expertise and background to participate in the discussion as a moderator and co-creator of knowledge *with* the students.

We now give a basic overview of the format and share some reflections on its benefits and shortcomings. The intended audience for this post are instructors planning to offer a paper-reading seminar. We have collectively tested this format four times in formal course offerings. We also expect that a similar role-playing format could be effective in informal reading groups or journal clubs.

3 The Roles

In past offerings (UNC 2020, UofT 2020, UofT 2019, Columbia 2015), Colin and Alec have offered roughly the same format on the topics of Deep Learning with Limited Labeled Data and Geometry Processing & Computer Animation, respectively. We have structured the course as follows: In each (bi-)weekly session, we discuss a single research paper, for which each student has been assigned a specific role. In our past offerings we have used the following roles, with accompanying instructions provided to students:

Scientific Peer Reviewer. The paper has not been published yet and is currently submitted to a top conference where you've been assigned as a peer reviewer. Complete a full review of the paper answering all prompts of the official review form of the top venue in this research area (e.g., *NeurIPS* for Deep Learning and *ACM SIGGRAPH* for Geometry & Animation). This includes recommending whether to accept or reject the paper.

Archaeologist. This paper was found buried under ground in the desert. You're an archeologist who must determine where this paper sits in the context of previous and

subsequent work. Find and report on one *older* paper cited within the current paper that substantially influenced the current paper and one *newer* paper that cites this current paper.

 **Academic Researcher.** You're a researcher who is working on a new project in this area. Propose an imaginary follow-up project *not just* based on the current but only possible due to the existence and success of the current paper.

 **Industry Practitioner.** You work at a company or organization developing an application or product of your choice (that has not already been suggested in a prior session). Bring a convincing pitch for why you should be paid to implement the method in the paper, and discuss at least one positive and negative impact of this application.

 **Hacker.** You're a hacker who needs a demo of this paper ASAP. Implement a small part or simplified version of the paper on a small dataset or toy problem. Prepare to share the core code of the algorithm to the class and demo your implementation. Do not simply download and run an existing implementation – though you are welcome to use (and give credit to) an existing implementation for "backbone" code.

 **Private Investigator.** You are a detective who needs to run a background check on one of the paper's authors. Where have they worked? What did they study? What previous projects might have led to working on this one? What motivated them to work on this project? Feel free to contact the authors, but remember to be courteous, polite, and on-topic.

 **Social Impact Assessor.** Identify how this paper self-assesses its (likely positive) impact on the world. Have any additional positive social impacts left out? What are possible negative social impacts that were overlooked or omitted?

With a course enrollment of 10-20 students, assigning multiple students to the same role in each class session works well. In particular, Alec found for this size and his 110 minute class duration, it worked well to assign one or two students as Hacker, two students to each of the Academic Researcher, Industry Researcher, Private Investigator or Social Impact Assessor roles, and all remaining students as Reviewers and Archaeologists.

If enrollment is significantly larger than the number of roles, it may not be feasible to fit contributions from each student into each class session. To address this in Colin's 26-student class, he had students team up in small groups for each of the roles. Students alternated whether they were presenting class-by-class. Those that were not presenting were required to complete a small assignment by coming up with an alternate paper title and submitting at least one question or observation about the paper. Then, after the presentations finished in the class, discussion was seeded based on points raised in the presentations and the submitted questions.

All students read the assigned paper before each class session. Students in different roles will prepare for their contribution to the discussion in different ways. For example, the Hacker spends most of the out-of-class time writing a software demo, while the Private Investigator is scouring the web and emailing with the paper's authors. In Alec's course, all students are asked every week to bring an alternative title or missing experiment. At the beginning of each class, these are used as a rapid-fire ice breaker. To keep this small weekly task interesting, constraints were sometimes added to inspire creativity, such as "*this week your missing title must alliterate*", which led to the following titles for the paper "As-Rigid-As-Possible Surface Modeling" by Sorkine & Alexa:

- An Alternating Algorithm for ARAP Alterations by Alexa and Another Author
- Cellular Changeless Cost Convergence
- Detail-Defending Deformation
- Efficient Energy Enabled Editing Enforcing Edge Elegance
- IEEE: Iterative Edge-Based Energy-Metric Editing
- Localized Laplacian LL^T for Life-Like Locomotion
- Manner Maintained Mesh Model Morphing Method
- Olga's Oddment Obscuring Oscillating Optimizations on Outer-shells
- Proper Property Preserved Process
- Repeatedly Reshaping Rigid-ish Recipients
- Robust Rigid Reshaping with Repeated RMS
- Sorkine's Scantily Stretching Structured Surface Shaping

Geometry processing as a subfield of computer graphics is very visual, so Alec encouraged students to use slides *optionally* for visual aids. To minimize time spent context switching or fighting with screen sharing/projector dongles, students were encouraged to organize any supplemental images, videos or slides into a shared Google Slide document. By requiring students to title their slides with [role emoji]: [student name] (as in 👩🏽‍🎨 : Sally), slides could be organized and quickly identified during the session. Students were not required or expected to prepare a full-blown presentation.

In physical offerings, students do not leave their seats but rather add their contribution in a round-table environment. Ideally, this is realized by physical arranging desks in a circle. In a virtual offering, the "gallery view" of Zoom or other video conferencing software works as a direct replacement.

The class period progresses from the ice-breaker in an *ad hoc* tour of the roles as the conversation dictates. Natural discussions emerge between students with the same role: for example, Reviewers may disagree on whether paper should be "accepted". If not, the instructor could push in the opposite direction of consensus to prompt further justification.

Discussion from one role often leads naturally into others. The project proposal from a PhD Student may hit on a desired application mentioned in one of the papers found by an Archeologist. It was not uncommon to hear students jump in with interjections like "Oh, in the paper from the future I read they …" The Hacker role, which may require a screen-sharing/projector disruption, provides a useful halfway point for the class time.

The workload for the instructor outside of class is minimal. They determine a candidate list of papers, e.g., before the term starts or *ad hoc* based on the natural direction and interests that evolves among the students. They should also read the paper each week. Assigning and ensuring a somewhat uniform permutation of students through roles is necessary each week (and admittedly a little tricky in the early weeks while enrollment may still be fluctuating).

So far, all of the courses in this format have been graduate-level. Consequently, assigning marks for student work has been uncomplicated (and admittedly generous). Each week students are marked according to simple thresholds for their participation. Double weight is placed on the week that a student plays the Hacker role (and they play this role exactly once per term) since this role requires more work. The instructor gives examples of how to simplify a paper for rapid implementation: For example, in Geometry Processing and Computer Animation most papers are for 3D shapes but their methods can be simplified to 2D. The instructor may consider Hacker-ability when choosing suitable papers for the list by ruling out papers that require elaborate system building or non-public datasets. The final term marks — the average of weekly marks — still present variation, albeit tending closer to 100%.

4 Reflections

Overall, we have found that students thoroughly enjoyed this course format. For example, in course evaluations students have said:

- I really enjoyed the class and taking it has given me more confidence when reading and implementing work from recent literature.
- The environment of the seminar was comfortable and conducive to discussion and learning.
- The different roles taken gave us new perspectives on how to critically evaluate research.
- The focus on self–guided learning was good and allowed for students to learn at their own pace.
- I really like how this course is organized! The role–playing part is an excellent way of learning different aspect of literature review and how to write a good paper.

Beyond this positive student feedback, we both have really enjoyed teaching classes in the role-playing seminar format. The standard one-to-many seminar format can devolve into little or no analysis: The presenter presents the paper as-is and the audience asks banal questions because maybe they just skimmed or didn't read the paper. We have been struck by how deep the analysis can go when students are role-playing. Hearing the students give critical takes on papers, come up with creative research ideas, and share their experience reimplementing the paper is much more interesting than a simple summary. It also leads to much better opportunities for discussion.

From the instructor's point of view, this format provides significantly more interaction with each individual student than the standard format. When writing letters of reference, we could draw from an entire term's worth of interactions rather than one or two presentations.

There are some things we would change or refine in future offerings of this course:

In Colin's offering, he wasn't sure what the interests and backgrounds of the students would be since it was his first class at UNC. Instead of coming up with a specific schedule of papers ahead of time, he created a long list of possible papers and had the students vote on which ones they wanted to read. The class covered the papers that received the most votes. This was a mistake because some papers ended up being redundant with one another. In the future, it would be better to curate the list (which is the approach Alec has taken).

It has been a bit difficult to get the "Social Impact Assessor" role right. In earlier iterations of this format, Alec defined the role as a "Ethicist from the Future" tasked with assessing whether the paper has made a positive or negative impact on the world. Unfortunately, most students described similar impacts (regardless of the paper's content) that related only to the general course topic (often describing unrealistic, apocalyptic Terminator-style scenarios). Colin chose instead to absorb impact assessment into the "Industry Practitioner" role, requiring articulation of negative and positive impacts. Unfortunately, this resulted in a lot of overlap between the "Practitioner" and "Researcher" roles, and most students did not engage meaningfully with

brainstorming negative impacts. In Alec's most recent offering, the role evolved to the "Social Impact Assessor" described above. Students were cautioned that this is a role computer scientists often have the least training and experience with. To prepare for this, students were encouraged to consult the recent blog post about impact assessment in computer science peer review [hecht2018]. For the first week of the class, *all* students were required to read a recent article about insidious racism in Computer Animation research [kim2020] as an example of critical analysis. To further strengthen this role, Alec had students read Invisible Women [criado-perez2019] over the course of the term and devoted the final class to discussing it.

Students in Colin's offering told him at the end that they would have valued hearing his perspective and opinion more often. He had framed the course so that it was a perhaps *too* egalitarian discussion, which maybe overly prevented the students from receiving his "expert" opinion on the papers.

The Hacker role requires a more serious time commitment. Students in Alec's offerings often requested knowing in advance which week and for which paper they would play the Hacker role so that they could plan in advance. Unfortunately, supporting this was at odds with *ad hoc* role assignment and fluidly determining the next paper based on current interests and discussions.

So far our graduate course offerings have been small relative to the undergraduate courses at public universities that often enroll hundreds of students. Based on the success of this format, we wonder how this format could be scaled up to large, undergraduate courses. Moreover, we wonder whether role-playing discussions as a form of active learning replace the classic lecture-assignment model entirely. That is, could an introductory or core computer science courses curriculum be supported fully or in part by in class role-playing exercises? We are keenly aware that our success inherits a strong possibility of selection bias as graduate students are typically more experienced than undergraduates and often select advanced courses according to their research interests. Furthermore, marks are far less critical in graduate programs focusing on research. More accurate assessment may be required for undergraduate deployment.

Upon writing this post, we discovered we are not the first to consider role-playing in paper-reading seminars. Parrott & Cherry [parrot2011] describe similar motivations to ours and share similar (albeit more formal) observational findings for the success of their five roles: Discussion Leader, Passage Master, Devil's Advocate, Creative Connector, and Reporter. In other contexts, explicit role assignment has been reported to increase engagement in online breakout rooms [mcmurtrie2020]. Meanwhile, rotating interactions between students could be linked to reducing racial conflict in primary education [aronson1978].

We are excited that instructors at other universities are already offering seminars inspired by our successful role-playing format (for example, by Ismini Lourentzou at Virginia Tech). We are eager to incorporate their experiences, feedback, and modifications.

References

[**freire1968**] Paulo Freire. *Pedagogy of the Oppressed*, 1968.

[**hecht2018**] Brent Hecht, Lauren Wilcox, Jeffrey P. Bigham, Johannes Schöning, Ehsan Hoque, Jason Ernst, Yonatan Bisk, Luigi De Russis, Lana Yarosh, Bushra Anjam, Danish Contractor, and Cathy Wu. "It's Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process." *ACM Future of Computing Blog*, https://acm-fca.org/2018/03/29/negativeimpacts/, 2018.

[**kim2020**] Theodore Kim. "The Racist Legacy of Computer-Generated Humans," *Scientific American*, 2020.

[**criado-perez2019**] Caroline Criado-Perez. *Invisible Women: Data Bias in a World Designed for Men*. 2019.

[**parrot2011**] Heather Macpherson Parrott and Elizabeth Cherry. "Using Structured Reading Groups to Facilitate Deep Learning," *Teaching Sociology*, 2011.

[**mcmurtrie2020**] Beth McMurtrie. "Teaching: How to Make Breakout Rooms Work Better," *The Chronicle of Higher Education*, 2020.

[**aronson1978**] Elliot Aronson. *The Jigsaw Classroom,* 1978.