Proceedings of the Seventh Symposium on Educational Advances in Artificial Intelligence (EAAI-17)

A Summer Research Experience in Robotics

Cindy M. Grimm, Alicia Lyman-Holt, William D. Smart
Oregon State University
Corvallis, OR 97331
United States

Email: {cindy.grimm,alicia.lyman-holt,bill.smart}@oregonstate.edu

Abstract

The Robotics Program at Oregon State University has been running an NSF-funded summer Research Experiences for Undergraduates (REU) site since 2014. Over twenty students per year (on average) have participated in the site, spending ten weeks embedded in the OSU Robotics Program. Our main focus with this REU Site is to give the participants a complete research experience, from problem definition to the final presentation of results, "in miniature". Our secondary educational objectives are: 1) Teach basic non-technical skills needed for graduate work, such as time management and literature review, 2) Provide details on how to apply to graduate school and for funding, 3) Clarify what we look for in a graduate student, and 4) Detail what to expect from the graduate student experience. In this paper, we describe the overall structure of the participants' summer experience, outline some of the training materials that we use, describe the motivations for our approach, and discuss the lessons that we have learned after running the program for a number of years.

Introduction

We have been running an NSF-funded summer Research Experiences for Undergraduates (REU) site, "Robots in the Real World"[1] in the Robotics Program at Oregon State University (OSU) since the summer of 2014. At the time of writing, after the third summer of the program, 65 students have participated — 34 supported by the site award, and a further 31 supported from other sources.

The Robotics group at OSU comprises eleven faculty with a variety of backgrounds from computer science to mechanical engineering to applied mathematics, and covers all aspects of robotics from mechanical design to human-robot interaction. The group occupies approximately 18,000 square feet of space in Graf Hall on the main OSU campus, comprising faculty and graduate student offices, social areas, and a large shared laboratory space (figure 1).

Undergraduates, graduate students, postdoctoral scholars, and faculty work in the same spaces, and researchers with backgrounds in mechanical engineering sit next to computer scientists and applied mathematicians. Our REU students are added into this mix, being treated just like any other member of the robotics research community for the duration of their summer experience.

In this paper, we describe the goals we have for the program and give a brief summary of how well we met them. We then describe the culture of the robotics group at OSU, the overall structure of the program, and how these explicitly interact with each other. We then detail the programmatic elements that we developed, both project-related and more general educational elements. Finally, we share some of the results from our formal evaluation of the program, and discuss some of the lessons we have learned from running the program for three years.

Program Goals

Our REU program was designed with multiple goals in mind. Perhaps the most important was to identify and recruit high-quality graduate students to the OSU Robotics graduate program. At the time we started the program in 2014, we had submitted the paperwork for our M.S. and Ph.D. programs in Robotics, and were trying to get the word out about our nascent programs. OSU has not, historically, been associated strongly with robotics, and we felt that a strong REU summer program was one way to change this. The REU program gave us an excuse to contact our peers across the country, asking them to tell their undergraduates about the summer program, and also letting them know about our graduate degree programs. Also, having students who have participated in the program return to their undergraduate institutions has allowed us to spread the word about OSU at the grass-roots level. To help with this, we have made sure that all of our participants have concrete products from their work here that they can show to their peers: posters, papers, videos, etc.

A secondary goal was to give our own graduate students some experience mentoring undergraduates who are not already part of the group. Often, undergraduate students who get involved in projects are already known quantities — they have been volunteering in the group, or have been in our classes. They already understand the culture of the group, and are likely familiar with the work we're doing. New students, from outside of OSU, are different, and require a different skill set to effectively mentor. By explicitly pairing up our REU students with graduate student mentors, we give the graduate students some experience of bringing someone entirely new into a project, supporting them, and making

[1] http://robotics.oregonstate.edu/reu

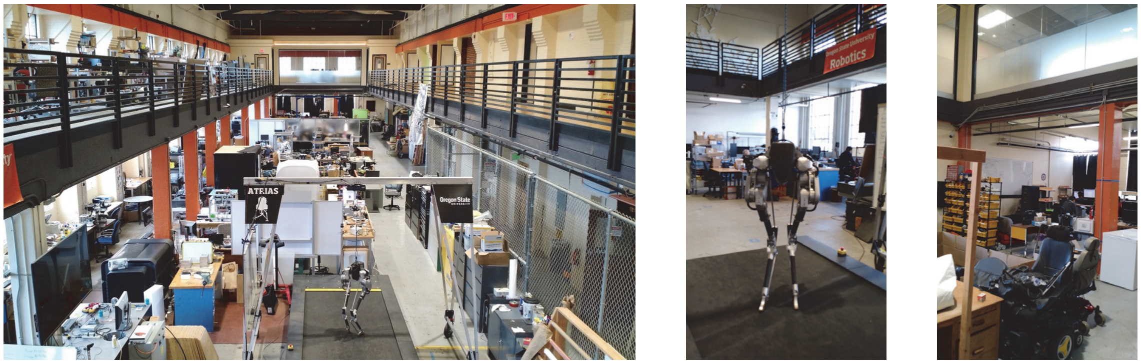

Figure 1: The shared laboratory space in Graf Hall, on the main OSU campus.

sure that they are productive in a very limited period of time. This reflects our belief that the program is a two-way street: the participants learn something from their experience, but the mentors also learn something equally as valuable.

REU Culture and Project Goals

Our goal is for REU students to complete projects that are meaningful and necessary within the context of a larger research agenda; the work may not be publishable in and of itself but it should contribute to a publication. It should **not** be a toy problem thought up just for the REU program. There are two key components to making this work: 1) A well-established hierarchical mentoring structure with strong cross-disciplinary ties, and 2) A broad list of subprojects suitable for novice researchers. We match projects to the skill sets of the students (eg, ensuring that a student has sufficient programming experience for largely coding projects) but have actually found that, beyond broad matching, this is not as important as establishing good mentoring. Students are remarkably adept at picking up new skills on the fly provided they feel supported and guided in the effort.

Culture, Mentoring, and our Physical Space

Most of our research groups are structured hierarchically: one or more faculty advisers, perhaps a post-doctoral scholar, one or more graduate students, and a small number of undergraduates. We expect our post-docs to mentor our graduate students, and our graduate students to be actively advising the undergraduates. There is also an expectation that students across all levels will help other students. Our lab space is organized by projects, rather than by group, with the boundaries intentionally fuzzy (both physically and intellectually). Graduate students from different areas share office space, and work with each other across project boundaries. Many projects have multiple faculty members involved, and again we expect that any student can go to any faculty member for help with problems in their area.

All incoming REU students are paired with existing graduate students and undergraduates and, if possible, with other incoming REU students [2]. All REU students provide feedback to each other during our weekly meetings. This not only reduces isolation, but also makes the students feel like they are part of a team, and increases the likelihood that they will stay on-track and produce a more substantial project. It also gives them a more complete experience of what research in robotics is actually like. We firmly believe that there are (very) few interesting single-person projects in robotics these days; collaboration skills are absolutely vital, even for graduate students specializing in a focused technical area.

Concrete suggestions: (1) Add REU students to an existing project, rather than giving them something isolated to work on. (2) Make sure that there's a dense network of peer- and near-peer-mentoring.

Project Specification

All of our REU projects are part of a larger, on-going research agenda or project. They fall roughly into three categories: 1) **Incorporating new functionality into an existing in-use system**, *Examples: sonar for an underwater robot, adding liquid metal printing to an in-house silicon printer, better tracking of usage and development statistics in the Robot Operating System (ROS) (Quigley et al. 2009)* 2) **Exploring alternatives to existing functionality**, *Examples: a different ankle actuator design, an alternative force-transfer design, trying a new version of a reinforcement algorithm* and 3) **Implementing a prototype where the design or approach is already fairly clearly defined**, *Examples: building a tethered puck, building an object-return system for grasping*. All of these example projects have clear boundaries, are feasible to accomplish within ten week, and are situated with a larger research agenda.

Concrete suggestions: (1) Begin with an existing research endeavor and identify a "wish list" based on the three categories listed above. (2) It is better to reduce the number of projects and increase the number of students on each project, particularly if the project has a limited number of available mentors. (3) Avoid projects that are all-or-nothing; there should be two or three achievable sub-goals (which could be either experiential or educational).

[2]Modeled after the Computing Research Association's Distributed Mentor Program.

Weekly Structure

We bring the students together once a week for a faculty-led, two-hour meeting. Meetings are a mix of passive lectures and active, hands-on activities to implement the material in those lectures. Our primary goal is to keep students "on-track" with their projects; secondary goals are to 1) Train them in basic skills needed for graduate work, 2) Educate them on how graduate life is different than undergraduate. We expect that the students will receive the bulk of their project-specific research guidance from their mentors so the materials we provide are not tied to any specific research area or discipline. In the past, we have successfully combined forces with other REU programs for some of these meetings, since they cover skills that are widely applicable.

The topics we covered in the 2016 version of the REU Site are shown in table 1. Each week students have a "homework" assignment, which is expected to take between 2-6 hours to complete, and helps them build a portfolio of documentation for their technical work. The weekly meetings address three inter-related threads: project process, graduate skills training, and graduate school (or industry) expectations, which we discuss in the sections below. For each deliverable we supply the students with a template to provide guidance on what we expect (in addition to information given in the weekly meeting). Slides and project templates are available upon request from the authors.

Project Process

One of our primary goals is to accurately replicate the research experience "in miniature" in the (short) ten weeks that our participants are with us. We focus on teaching the students how to articulate their project's goals in the context of the overall project, and how to self-evaluate progress in light of those goals. One key point we repeat over and over is that, unlike a homework assignment, there is no "right" answer; part of defining a research project is defining success. We reinforce project goal articulation with three lectures (weeks 1, 4, and 8), three student project presentations (initial, mid, and final), three written artifacts (proposal, banner and research posters), and three small-group evaluation activities (evaluating project time lines, elevator pitches, and poster draft). In addition, many advisers will ask students to write the draft of a conference or journal paper based on their work over the summer.

Initial project presentations: Our initial lecture focuses on why they are doing project proposals and the elements that go into them. We explicitly break out project goals, implementation, and evaluation — asking them to create a time-line (evaluated in week 3), identify potential bottlenecks and alternative approaches. The templates provide very specific guidance on what to write and how much (one template for the initial project write-up and one for the presentation). This information is given on the first day, with specific instructions on eliciting this information from their mentor.

Students present their projects to the entire group and some subset of mentors at the start of week 2. To help the students critique each other we ask them to imagine that they would be asked to step into the project — actually implement it. We find that this prompt helps students to identify missing or confusing information in the presentation. We also provide meta-presentation comments (e.g., on timing, images/graphics, organization).This is also where we identify students who's projects are ill-defined.

The **mid-term presentation** (around 5-10 minutes) is a status update and consists of three questions: Re-iterate the goals, present progress so far, and explain any stumbling blocks or on-going struggles. If necessary, we may ask the student to re-define their project goals or approach at this point. The final presentation is longer (10-20 minutes) and is a standard, short research talk.

Our objective is for students to be able to clearly state the project goals and separate these from the project *implementation*. Students give some form of project goal statement a minimum of 6 times (3 oral, 3 written). This reiteration both reminds the students of what they're trying to do and reinforces presentation skills (how to present your work succinctly to a general audience). Explicitly separating the problem from the implementation —- and asking them to give alternative approaches — keeps students focused on what they're trying to accomplish (as opposed to getting lost in the how).

Students provide mid-term goals and a time-line in the initial presentation, along with fall-back approaches. We explicitly encourage students to list what they've tried and what worked (and what didn't) in their mid-term presentation in order to emphasize process over the result. This is particularly useful for projects that turned out to be harder than expected, and also serves as a calibration mechanism (did things take as long as you expected?).

We also ask students to include an evaluation mechanism, and emphasize the importance of knowing what "success" looks like and how to measure it objectively. This also helps them to think about how to measure their own progress.

Concrete suggestions: (1) Focus on teaching that research is a *process*, not simply an end-goal. (2) Remind them (more than once) that failure is all part of the process. (3) Provide very specific instructions (and examples if possible) on what you want the students to do. (4) Provide multiple opportunities for students to critique each other's work.

Graduate Skills Training

The skills we teach, in addition to the research skills outlined in the previous section, are: 1) literature review, 2) critiquing, and 3) documenting. We focus on these three elements because they are often neglected, or students are expected just to pick them up.

Literature review: We show the students the Mendeley citation management system, and talk about how to read a paper (get the big idea, look at the results, skim where possible, and chasing references). How many papers a student reviews as part of their technical work depends on the project: somewhere between 5 and 20 is typical.

Critiquing: We include a critiquing lecture in week 5, and before every hands-on session we provide guidance on how to provide *structured* feedback for that session. In the critique lecture we focus on how to structure a critique to get the most out of it (moving from positive aspects to potential problems to suggestions), critique language (how to present

Wk	Topic	Activity	Assignment
1	Defining a research project	Introductions, logistics	Project presentation
2	(No lecture)	Project presentations	Project timeline
3	Grad student skills and tools	SG Timeline evaluation	Project "elevator pitch"
4	Time management	Project "elevator pitches"	Banner posters
5	Critiquing technical work	SG Banner poster critique	Mid-project presentations
6	Documenting your work	Mid-project presentations	Documentation draft
7	Grad school, NSF GRFP application process	SG Documentation review	Application rough draft
8	Designing a research poster	SG Grad student application	Poster draft
9	Picking grad school, adviser Building professional network	SG Poster review	Final presentation
10	(No lecture)	Final presentations	

Table 1: Activity and weekly meeting time line. SG indicates small-group, round-robin style peer review. Banner posters are simplified research posters that summarize the project goals and approach (suitable for viewing on a banner from a distance).

your suggestions so they'll be listened to) and *how* to take a critique (focus on the content). Our approach is based on a combination of experience with art-based critique and how professional facilitators guide brainstorming sessions.

Documentation: Students document the materials and processes developed for their project (specific format depends on their research group's procedures). We focus on what should be in documentation and do a hands-on activity to critique each student's documentation to ensure key questions are answered. By this time we usually have multiple groups that have run into poor (or non-existent) documentation, which helps to reinforce this skill's importance.

For every small-group activity we provide a list of questions or types of questions, then let the students self-form into small groups with a volunteer in each group to start the process. The faculty mentors move from group to group, first asking prompting questions and, only after the students have begun self-critiquing, offering suggestions themselves.

Concrete suggestions: (1) Take the lessons you've learned over your career, and distill them down for the students, so that they don't have to repeat your mistakes. Even if the lessons seem obvious to you now. (2) Emphasize to the students that these things will help them, and they should try them out, even if they're skeptical.

Graduate School Expectations

The third thread in our weekly meetings is the expectations that the students should have when they enter graduate school. This is mostly informative, although we do interactively critique student's graduate school or job application cover letters. We discuss the application process — how are graduate students selected? — and how graduate student life differs from undergraduate. Our goal is partly to demystify the process, but also to help students set priorities when they return to their institutions, so that they can focus on the things that will materially help them get into the graduate program of their choice. We discuss the relative importance of the different elements of the application (GPA, GRE scores, application letter, research projects), what to put in the application letter, how funding works (GRA versus GTA, relationship of funding to project, applying for grants), how to select an adviser or university, the difference between a Master's and a PhD (both getting the degree, and what you can do with it afterwards), the variables that go into making offers (available funding, skill set match), and how to make contact with a potential adviser (mostly, what *not* to do). We also help the students create a Linked In profile, and talk about the importance of building a professional presence on the web.

Concrete suggestions: (1) Much of the graduate school process is opaque. Exposing *your* approach and evaluation criteria is a great way to shed some light on it. (2) Be honest with the students about what *actually* matters when you evaluate graduate student applications. Is a 3.9 GPA actually better than 3.8, or does a good reference letter trump everything?

Evaluation

Figure 2 shows the geographical distribution of our participants. Note that OSU is located in Corvallis, OR, which corresponds to the leftmost yellow marker at the top of the map. We are pleased to note that we have successfully recruited participants from all parts of the United States. We received 77 applications in year one, 266 in year two, and 155 in year three. The low number of applications in year one is due to the late award date of the proposal, and subsequent late start to recruitment. This makes our aggregate acceptance rate over the first three years of the program 6.8%. Figure 2 shows the self-reported ethnicity and gender of the participants. While the numbers do not quite reflect the general population of the United States, we believe that they are a good start for a technical field like robotics. The demographics in our applicant pool show a similar composition.

Our program has been both formally evaluated through an anonymous survey (primarily on how well we communicate the graduate experience material described above), semi-formally through structured interviews (mostly focused on process and logistics), and informally by simply asking the students what worked and what didn't after interactive sessions. Our formal evaluation survey consisted of eight 7-point Likert-scale statements, given before and after the

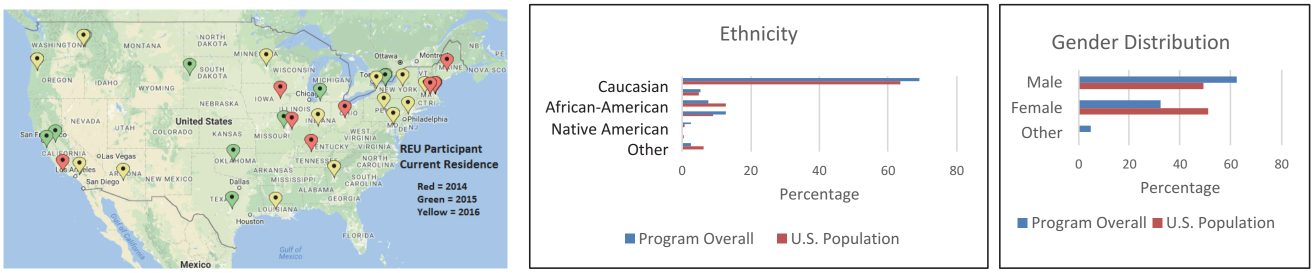

Figure 2: Left: Current residence locations of REU participants, at the time they participated in the program, for students supported by the REU site award (home institution). Middle: Self-reported ethnicity, by year, compared to 2010 census data. Right: Self-reported gender by year.

summer, described below and shown in figure 3. None of the results show statistical significance (t-test at 5% level), due to a small sample size, although they do show a clear trend.

Question 1: I understand the differences between undergraduate and graduate studies. Overall, students appear to enter the program with a fairly high self-reported understanding of the differences between undergraduate and graduate students. In all years of the program, there was a reported increase in this understanding, which amounted to a 17.5% change, a 13.5% change, and a 13.05% change, respectively.

Question 2: I understand the coursework requirement for graduate studies. The students in the program showed an increase in their self-reported understanding of the coursework requirement for graduate studies in both years of the program. In Year 1 the increase was 39.4%, in Year 2 the increase was 25%, and in Year 3 the increase was 44.72%.

Question 3: I understand the research expectations of graduate studies. The students in the program showed an increase in their self-reported understanding of the research requirement for graduate studies in both years of the program. In Year 1 the increase was 31.4%, in Year 2 the increase was 19.4%, and in Year 3 the increase was 34.9%.

Question 4: I understand how graduate students are typically funded. In general, students reported an increase in their understanding of how graduate students are funded.

Question 5: I understand how to apply for external funding. In general, students reported an increase in their understanding of how to apply for external funding.

Question 6: I am confident that I can write a successful application for external funding. Students reported an increase in their confidence to write a successful application for external funding.

Question 7: I plan on applying to graduate school. There did not appear to much of an increase in students' self-reported plan to apply to graduate school.

Question 8: I am confident that I can get into the graduate program that I want to. In Year 1, there was an increase in students? self-reported confidence in being accepted into the graduate program of their choice. In Year 2 there was no difference in students? self-reported confidence, and in Year 3 there was a small increase in students? self-reported confidence.

Almost all of the questions show an improvement after the program, for each of the years of the program and, based on this, we believe that we have been successful in the goals for this previous award.

As indirect evaluation, several of the projects from the first two summers have been published (Hubicki et al. 2016; Bowen-Biggs et al. 2016; John et al. 2016; Hubers et al. 2015a; 2015b; Holloway et al. 2015; West et al. 2016; Cesare et al. 2015; Curran et al. 2015; Faraji et al. 2015; Allani et al. 2016; Faraji et al. 2015; Sundberg et al. 2016), with the ones from this summer in preparation.

Discussion and Lessons Learned

After three years of running an REU site focused on robotics, we have learned a few things. Some of these are, in hindsight, somewhat obvious. Others are less so.

It's a surprising amount of work to run one of these, and it really, really helps to have an administrator to share some of the burden. Many of the tasks, such as refreshing the content on the web site and getting the applications in place, are very time-sensitive. Missing the (self-imposed) deadlines and doing these things a week or two late will materially affect the number of students who apply to your program.

It's a good idea to **coordinate your dates with other programs around the country**, since many students will be applying to several programs. We've lost more than one good applicant in the past because we could not give them a definite decision before the deadline from another program.

Students do much better when they are **working on a small part of a larger project** rather than on a small, standalong project. They integrate better into the group, are able to see their work in a larger context, and get more involved in the work. At the same time, it's important to **carefully draw boundaries around each project** so that the students can take unambiguous ownership of the work.

The important word in Research Experiences for Under-

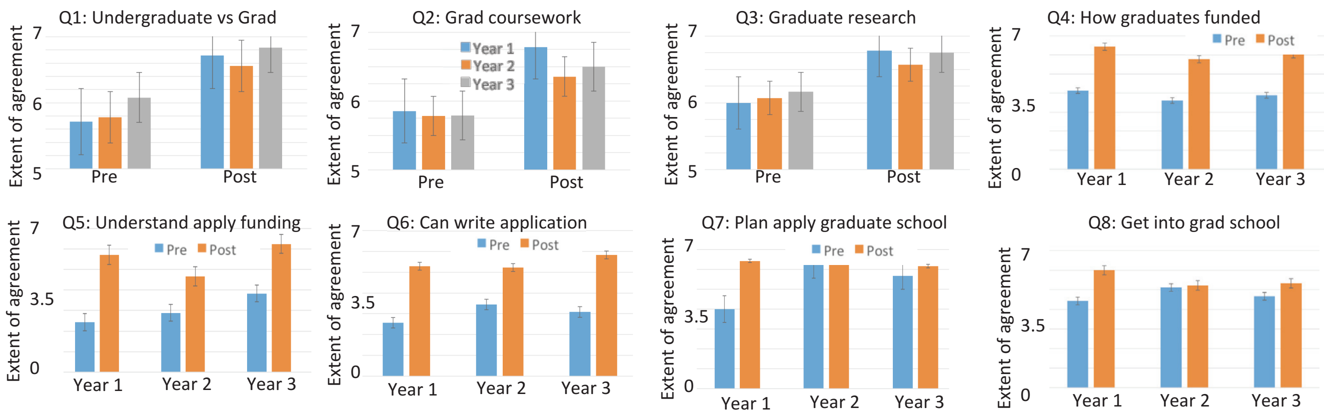

Figure 3: Results from the eight evaluation questions.

graduates is **Experiences**. An REU program is not a cheap source of research labor over the summer, it's an investment in the future; either your own group or the community as a whole. Ten weeks is a really short time to do anything interesting, especially if the students have limited prior research experience. If you could go from idea to publishable paper in ten weeks, then we'd all be publishing five papers a year as a single author. It's important to remind the students that research is a process built of many failures and, should they get to the end of the summer without publishable work that's fine, as long as they (personally) learned something and also pushed the boundaries of knowledge back a bit. If all that happens is a particular student knows with more certainty whether or not they want to go to graduate school, that's a successful outcome.

And, finally, as we've long suspected, **robotics is a great focus for an REU site**, since it attracts a diverse collection of students who are interested, demographically and in terms of academic discipline.

Acknowledgments

The NSF REU Site described in this paper was supported by the National Science Foundation, under award CNS 1359480.

References

Allani, S. K.; John, B.; Ruiz, J.; Dixit, S.; Carter, J.; Grimm, C.; and balasubramanian, R. 2016. Evaluating human gaze patterns during grasping tasks: Robot versus human hand. In *SAP*. ACM.

Bowen-Biggs, L.; Dazo, S.; Zhang, Y.; Hubers, A.; Rueben, M.; Sowell, R.; Smart, W. D.; and Grimm, C. M. 2016. A method for establishing correspondences between hand-drawn and sensor-generated maps. In *ICSR*.

Cesare, K.; Skeele, R.; Yoo, S.-H.; Zhang, Y.; and Hollinger, G. A. 2015. Multi-uav exploration with limited communication and battery. In *ICRA*, 2230–2235. IEEE.

Curran, W.; Thornton, T.; Arvey, B.; and Smart, W. D. 2015. Evaluating impact in the ros ecosystem. In *ICRA*, 6213–6219.

Faraji, H.; Veile, S.; Hemleben, S.; Zaytsev, P.; Wright, J.; Luchsinger, H.; and Hatton, R. L. 2015. Impulse redirection of a tethered projectile. In *ASME 2015 Dynamic Systems and Control Conference*.

Holloway, M.; Sanandaji, A.; Yates, D.; Krigger, A.; Sowell, R.; West, R.; and Grimm, C. 2015. *ISVC*. Springer International Publishing. chapter Guided Structure-Aligned Segmentation of Volumetric Data, 307–317.

Hubers, A.; Andrulis, E.; Scott, L.; Stirrat, T.; Zhang, R.; Rueben, M.; Grimm, C. M.; and Smart, W. D. 2015a. Using video manipulation to protect privacy in remote presence systems. In *ICSR*, volume 9388, 245–254.

Hubers, A.; Andrulis, E.; Smart, W. D.; Scott, L.; Stirrat, T.; Tran, D.; Zhang, R.; Sowell, R.; and Grimm, C. 2015b. Video manipulation techniques for the protection of privacy in remote presence systems. In *HRI Extended Abstracts*, 59–60. ACM.

Hubicki, C.; Abate, A.; Clary, P.; Rezazadeh, S.; Jones, M.; Peekema, A.; J, J. V. W.; Domres, R.; Wu, A.; Martin, W.; Geyer, H.; and Hurst, J. 2016. Walking and running with passive compliance: Lessons from engineering a live demonstration of the atrias biped. *IEEE Robotics and Automation Magazine*.

John, B.; Carter, J.; Ruiz, J.; Dixit, S.; Allani, S. K.; Grimm, C. M.; and Balasubramanian, R. 2016. Human-planned robotic grasp ranges: Capture and validation. In *AAAI-HRI*.

Quigley, M.; Conley, K.; Gerkey, B.; Faust, J.; Foote, T.; Leibs, J.; Berger, E.; Wheeler, R.; and Ng, A. 2009. ROS: An open-source robot operating system. In *ICRA*.

Sundberg, M.; Litwinczyk, W.; Grimm, C.; and Balasubramanian, R. 2016. Visual cues used to evaluate grasps from images. In *ICRA*, 1965–1971.

West, R.; Kajihara, M.; Parola, M.; Hays, K.; Hillard, L.; Carlew, A.; Deutsch, J.; Lane, B.; Holloway, M.; John, B.; Sanandaji, A.; and Grimm, C. 2016. Eliciting tacit expertise in 3d volume segmentation. In *VICINI*. ACM.