

TALES OF A GROUP LEADER: BUILDING LEADING ACADEMIC RESEARCH GROUPS

Ravi Ramamoorthi; University of California, San Diego; October 30, 2021

1. PRELUDE

In late July 2006, I traveled to Boston for the annual SIGGRAPH Conference, the largest annual academic and industrial convention in computer graphics, attended by tens of thousands of participants, including thousands of researchers, from senior professors to industrial scientists to new students in the field. This happened as a deadly heat wave enveloped the North American continent in the summer of 2006, an early manifestation of the climate change catastrophes the world would soon reckon with regularly. However, it wasn't the searing triple-digit outside temperatures that were my immediate concern, but the heat of a controversy within the usually staid research community.

Oblivious to the heat outside, I had just arrived from New York (I was in the fourth year of my first faculty job as an assistant professor at Columbia University) and entered the air-conditioned confines of the Boston Convention Center to pick up my registration documents. As I walked towards the registration, I ran into Professor Julie Dorsey, a faculty member at Yale who was the papers chair presiding over the selection of technical academic papers for SIGGRAPH that year. Julie and her technical papers committee (which I had been invited to join for the first time that year) had selected 86 papers of the 474 submissions after careful review, and in a process that excluded committee members from papers they may be conflicted with as coming from their own groups, their institution or collaborators.

Julie's first words to me after a basic exchange of pleasantries were essentially "Congratulations on all your papers. I do think Columbia's papers were because of the quality of the work, rather than any other circumstances." The first part of her message was a gracious and well-accepted congratulatory note by a senior colleague and the papers chair. Indeed, my own research group and Columbia University had had a banner year at SIGGRAPH 2006, with 10 papers accepted for the conference, fully 12% of the 86 selected that year. I myself had 5 papers at the venue, while my erstwhile mentor at Columbia, Professor Shree Nayar was an author on 6 papers (two of which we collaborated on). But what did Julie mean by the rather cryptic second half of her message? What were the other circumstances she was talking about?

In subsequent discussions with colleagues and surfing the world-wide-web, I first learned of the controversy swirling around the field (admittedly, I must have been under a rock to not have known about it earlier). Academic publishing and academic reviewing are generally a demure enterprise, carefully searching for the truth, new ideas and new technology. While numerous books on science have indicated that scientists and researchers have human traits and human jealousies with grudges and enmities between scientists, these are rarely on public display. Moreover, in a new field like computer graphics, my experience had always been one of encouragement, with more senior researchers in the community encouraging the next generation, while keeping the standards of the reviewing process high as always, but fair and focused on the technical merits of the work.

This staid world had been disrupted a couple of months earlier by a post by Professor Michael Ashikhmin, then a junior faculty member at Stony Brook University. He had publicly announced that he was departing the field of academic computer graphics, not immediately as a result of tenure denial or dismissal, but out of frustration with academic reviewing of papers in computer graphics, and especially the SIGGRAPH review process. Needless to say, he had many of his papers rejected, in his mind unfairly. Bear in mind that SIGGRAPH is the premier conference in the field, and historically has had among the highest review standards in computer science, so it has always been difficult to get a paper into SIGGRAPH. These high standards were important historically to develop the respectability of computer graphics as an academic field, a respect that was hard to earn in the early days. Nevertheless, academic publishing is usually a requirement for advancement, tenure and awards, and the inability to get papers into SIGGRAPH can (and was for many junior faculty) a death-warrant for an academic computer graphics career.

Ashikhmin's post had gotten enough attention with the broader research community to create a blog with an active presence from a number of researchers in the field, and many more reading the information without participating. For

the most part, senior researchers felt the concerns expressed here, while having some merit with issues that should be addressed, were mostly overstated or devoid of clear factual basis, and the SIGGRAPH review process remained one of the gold standards within the academic community. Indeed, at a townhall in response later in the conference, SIGGRAPH 2007 papers chair Marc Levoy of Stanford university would forcefully assert that “The SIGGRAPH process is not badly broken.”

In any event, amongst the many messages and posts on the blog, was a comment wondering how Columbia had 10 papers at SIGGRAPH that year, when we had no previous history of being very active in the SIGGRAPH community, and how we had significantly more papers this year than any other university, even compared to traditional giants like Stanford (where I had received my Ph.D. in 2002). The post wondered aloud if it had anything to do with having us (myself and Prof. Nayar) on the SIGGRAPH papers committee for the first time. Of course, this claim is patently false, since papers committee members are not allowed to participate in any way on papers that involve their own institutions, and are excused from the room on papers they are conflicted on (it was indeed pleasant for me in my first time on the papers committee to leave the room so often on papers discussed in the top group at the start of the meeting). And it was this comment that Julie was referring to (and rejecting) in the second part of her message.

In many ways, I took this as an unusual mark of recognition and accomplishment. In other fields, one is often told to aim at not tepid growth but extraordinary success. One’s goal in starting a company should be not just to obtain a significant market share, but be so successful that the Feds hit you with an antitrust lawsuit. One’s goal as an athlete should be not just to produce better times, but break world records by so much that there are widespread (but untrue) allegations of drug-use. In contrast to these very public pursuits, it is not every-day in the private academic ivory tower that one is considered to be achieving success so significant that insinuations of unfair practices ensue. But when it does happen, it must herald an extraordinary level of achievement. And the recognition that followed was swift. Within the next two years, my senior Columbia colleague Shree Nayar became one of the younger members elected to the National Academy of Engineering. I myself received the SIGGRAPH Significant New Researcher Award for the top young researcher worldwide in computer graphics, in competition against the strongest field ever in SIGGRAPH history, and the Presidential Early Career Award for Scientists and Engineers at the White House, the highest early career honor from the US Federal Government, one of only 3 graphics faculty ever to do so to date.

Clearly, this wasn’t achieved by the “other circumstances” Julie referred to and rejected. But, how was it achieved? How did Columbia go from 0 papers at SIGGRAPH for the four years prior to my arrival in 2002, to be the leading computer graphics research group in the country and world four years later in 2006? This article discusses my role and insights in building world-leading computer graphics and vision research groups largely from scratch at Columbia, the University of California, Berkeley and the University of California, San Diego over a period of nearly two decades. It is intended as a manual or advice for faculty who may be leading or participating in multi-faculty academic research groups. While the intended audience is “group leaders” in science and engineering, the insights here are likely to be more broadly applicable, including in the context of industrial research labs, and junior faculty focusing more on their own careers rather than building up the larger group. There are many books and advice columns for a junior faculty member with respect to building their own research career, and I do not seek to duplicate this advice (though in many cases, I will touch on my own experiences and advice). However, it is increasingly the case that a department is measured by the strength of teams of faculty in particular research areas, such as computer graphics or computer vision in my own case (often referred together as visual computing), or in other areas such as theory, computer architecture, artificial intelligence, systems or security in computer science. Even if no official “group” or “group leader” is designated for these research areas, the faculty within them must band together to create an effective and cohesive research group to be most successful. At the other extreme, I have also seen and heard stories of dysfunctional research groups, which can be a real strain for all concerned. As such, these insights are intended for any faculty member involved in such a broader research group, and may also be of interest for others such as graduate students or even those outside the academy, to get an insight into what happens behind the scenes in academia. The events described here are based on my own recollection, and name other relevant participating faculty to be true to the situations as they happened. I apologize for any mischaracterizations and any unintended slights towards colleagues, most of whom are intended to be mentioned in the most positive light as exemplars of our community.

As I develop this article, I cannot emphasize enough the importance of mentoring. And I must acknowledge my own debt to my mentors, starting with Al Barr who agreed to take me on as a research student in my freshman year as an undergraduate at Caltech. Throughout my college and graduate school days and beyond, beyond the specific research insights, I was always overwhelmed by how deeply he had thought about issues in building a group not just for his

own research but also for the broader graphics group, from the placement of the fridge to the arrangement of different faculty's labs. I owe most of my intellectual development as a student first to the late great Jim Arvo, my senior/masters advisor at Caltech, and later to my Ph.D. advisor Pat Hanrahan, a once in a generation or once in a century faculty member who demonstrated for me what every aspect of intellectual and faculty life can be like at the highest levels (including building the greatest multi-faculty graphics groups ever created at Stanford). I did not start out with the aim of building multi-faculty research groups, but simply to be a great researcher in my own right as a student and later a faculty member. And if my first faculty search as a new Ph.D. from Stanford in 2002 had landed me in an established graphics group, that's where this story may have ended. But as fortune would have it, I landed up at Columbia, a university and group with great potential but no recent track record in computer graphics or SIGGRAPH. I owe much of my philosophy of group building to Shree Nayar, who not just mentored me as an individual faculty member but let me observe and participate in one of the most successful exercises ever to build a new group and center. In other circumstances, I might have gratefully accepted tenure at Columbia, and this story might have stopped shortly after the episode related at the start of this chapter. However, for better or worse, personal circumstances caused me to move twice, first to UC Berkeley where I had the great fortune to be mentored by Jitendra Malik, and finally to UC San Diego where I was recruited by Dean Albert Pisano to fully lead and build a new group. In this way, I've had the good fortune to see the group building effort from the ground up as a junior faculty member under a more senior colleague, as a newly tenured professor working with others at a similar career stage, and finally as the senior faculty member building a completely new group. Besides the mentors above, I owe thanks to all of my colleagues and many junior and senior faculty in the field who have provided advice and encouragement at various points. Finally, I owe an inordinate debt to my family and especially my wife Padmini Rangamani for all her support, encouragement and counsel over the past 18 years through essentially all of the adventures related here, and for requesting me to put down my insights in the written word. For those starting the incredibly exciting and fascinating opportunity to build a new research program, I hope these pages make your life that much easier.

2. FACULTY HIRING

Entering graduate school and becoming a new Ph.D. student is often described as a “sink-or-swim” experience, akin to being taken into the deep sea and left there. You either find your way, publish research and write a thesis (swim) or drown in the complexities of discovering new knowledge (sink). At the next stage, my undergraduate and master’s thesis advisor James Arvo, once likened being a junior faculty member to having to first write the grant proposals to finance a boat (and hire the crew, your team), so you could get back to shore. Within this paradigm, being a lab or center director involves doing all of the above, but is perhaps more akin to the early voyages of exploration; you need to outfit not just one boat, but an entire expedition into the unknown. Your first task is to hire the individual ship captains, or returning to the academic parlance, the faculty within your group or center. A similar role is played by lab directors in industry, who typically have a given headcount to hire junior and senior researchers. This chapter is also of interest to junior faculty and graduate students, since in any expansion of your faculty group, you will have an important role to play. Also note that many universities don’t have the luxury of hiring a senior faculty director in a new research area, and so you may be called to play many of the roles discussed here even as a junior faculty member.

Depending on what capacity you were hired into, and the timing, the current state of the group and existing faculty will differ, i.e., are you the first hire in a new research area the university is investing in, or are you the last of a sequence of hires? As such, the amount of hiring you can (or are able to) do varies. Regardless, in academia, each individual faculty hire is precious and must be used wisely, since faculty slots are usually very hard to find. From my own experience, I entered Columbia in August 2002 as the first of two mainstream computer graphics hires, but the third member of a broader vision and graphics concentration (with Shree Nayar and Peter Belhumeur), and a part of a broader envelope of 7 faculty within the Columbia Vision and Graphics Center (CVGC) formed by faculty with similar interests in January 2003. Building Columbia into a major player in computer graphics involved another key faculty hire, of Eitan Grinspun in 2004 (now at the University of Toronto, after 15 years at Columbia). It was this faculty nucleus, and the students we recruited and the research they did, that propelled us from having no presence in computer graphics, to becoming the leading researchers in four short years, as discussed in the previous chapter.

When I moved to UC Berkeley in January 2009, I was essentially the last part of a re-building and expansion of the computer graphics group (primarily with James O’Brien and Maneesh Agrawala [now at Stanford], who were at similar career stages then at the cusp of or just recently receiving tenure). Along with Jonathan Shewchuk and new human-computer interaction hire, Bjoern Hartmann, we formed the Visual Computing Lab (VCL) in July 2011. Being at the tail end of the faculty hiring effort was fortuitous, given that the great recession essentially froze new faculty slots for several years. When the economy improved and the faculty freeze thawed, I played an important role in hiring Alyosha Efros and Ren Ng (ultimately joined after I had departed). With this added strength, the year before I left in 2013-2014, I believe UC Berkeley and the VCL was the top group in the country in both computer graphics and computer vision.

With these experiences, I was ready for (and in any case was thrust into, ready or not) the challenge of a “clean-slate” leadership role within the Computer Science and Engineering Department at UCSD, of (re-)building the vision and graphics group from an essentially blank slate as a senior faculty group leader or center director. I am reminded of the experience Stefano Soatto from UCLA recounted to me more than a decade ago. When he was first a faculty member in Italy, on his first day, he arrived in an office that had just a chair, desk and a single notepad. When he joined UCLA, the situation was even more extreme; he entered his office the first day to find it completely empty. This is the blank, empty or clean slate that every new faculty member experiences in one way or another as an extraordinary challenge and an extraordinary opportunity. You are now “independent”, no longer guided or constrained by your Ph.D. or postdoctoral advisor. You start with nothing, except boundless energy, excitement and ideas to the brim. The future is what you make of it. Multiply this by an order of magnitude as a group leader or lab director starting a new research area. In addition to building up your own lab and research program, you must build the larger faculty group within your area. This chapter discusses my philosophies and advice on faculty hiring, motivated primarily by my efforts in putting together an extraordinary group of junior faculty within CSE at UCSD, and also drawing from my earlier experiences and those of others in the field.

Obviously, whom you hire as faculty colleagues in the area depends on opportunity, depending on who is on the market, who the department and you yourself are willing to get behind, and who is willing to come, often taking a risk in terms of joining a new faculty (as opposed to better established) group. Let’s pause for a moment though and think about the ideal sequence of building up the faculty group (indeed, at a typical universities, the competition for faculty

slots in the research area is tight, so the building up will happen over a period of many years... it took us 7 years at UCSD after my arrival to hire 4 junior faculty in CSE, and get to a critical mass of faculty). My first priority is to establish a nucleus of extreme strength in a particular research (sub)-area, usually my own. Indeed, I was hired at Columbia as part of a sub-group of faculty (Shree Nayar, Peter Belhumeur and myself) with similar interests in lighting and appearance; I broadened that expertise from physics-based computer vision to rendering in computer graphics. One of my attractions to UCSD was the similar strength it had in my primary research areas of computer graphics rendering and physics-based computer vision, between Profs. Henrik Wann Jensen and David Kriegman. With Kriegman's long leave to industry, the highest priority faculty hire I needed to make was a leading junior researcher and best paper award winner in physics-based vision, Manmohan Chandraker (now tenured) and subsequently upon Henrik's retirement, a leading new Ph.D. in rendering and SIGGRAPH Dissertation award winner, Tzu-Mao Li.. This ensured the (re-)building of a very strong nucleus around the areas of rendering, appearance and physics-based vision. With additional collaborations at UCSD and the local area, and unique equipment and facilities, these hires clearly established the university as a leader in this important research sub-area.

Why follow this approach, as opposed to first hiring faculty in a range of sub-areas to provide breadth? Part of this is admittedly self-serving; one wants to create a better (ideal) situation for oneself within the university by hiring faculty with which one will have very close collaborations, can co-advise and co-fund students, and work well together on educational and administrative matters. But it also has many practical benefits. You can't start from zero and immediately be the strongest group in the country across your entire research field. However, you can much more easily be the strongest group or one of the strongest groups within a more limited sub-area. Being a leader in the field has many benefits. On the one hand, the collaborative potentials, closer mentoring opportunities, and area leadership will enable you to hire and retain the strongest faculty in the area (including yourself), who might otherwise depart for higher-ranked universities or industry. On the other hand, there may be opportunities open to hire star researchers who may be overlooked at other universities that do not know or appreciate this sub-area to the same extent. Indeed, I view the four largest factors in one's success and hence choosing a university department as being the quality of collaborators, the quality of students, the quality of mentoring, and the availability of funding. Assume to a very crude approximation that the quality of students is determined by overall department and university ranking, that you don't really influence, and that funding is largely the same everywhere. Then, the primary factors you can influence are the quality of collaborators and the benefits you yourself will provide as a mentor. Moreover, having a strong cohesive (sub-)group will also enable you to recruit stronger students in competition with top universities. Ultimately, my goal as a faculty director or group leader is to make the current situation the ideal situation for both myself and other members of the group, so we can retain them, and they can be the most productive.

When I first plunged into faculty hiring at UCSD, I had occasion to receive some sage advice from Jessica Hodgins at Carnegie Mellon University, a winner of the Coons award, the highest lifetime achievement honor from SIGGRAPH, and someone who was hired in a similar role at CMU in 2000 to rebuild a graphics group from scratch. She proceeded to do this several times very successfully, and also build a research lab for Disney and subsequently Facebook. Her advice, quoted below, echoes many of the points above and also observes that it is an added bonus if the first faculty being hired have collaborative links to other parts of the department or university. *"Here are a few things that I thought that I did well at CMU when I was re-building the group (the first time). I looked for people that I thought could bridge to other parts of the school/department (Computer vision and robotics primarily). That allowed them to fit in better and get better students when the group was still struggling. I also looked for people that I thought could collaborate in various combinations as that made for a stronger group faster than strong individual contributors might have. Good luck -- it is always fun to have slots to spend!"*

Once one has built up the nucleus, one needs well-chosen faculty hires to expand into other sub-areas, in order to have more complete coverage of the field. My goal at UCSD was to hire a faculty member in computer vision and learning (ultimately filled by Hao Su), and one in graphics outside of rendering (ultimately filled by Albert Chern, a similar role for which I hired Eitan Grinspun at Columbia nearly two decades ago... coincidentally or otherwise, Albert and Eitan have the same Ph.D. advisor!). Often, the faculty involved will have similar intellectual philosophies, even if they are in different research sub-areas, which enables future collaborations. It's obviously harder to make the case about collaborations and mentoring than to the nucleus, but you can be an effective mentor even without working in exactly the same sub-field. Once the nucleus is established, the overall group should be a powerful recruiting pitch. By necessity or design, it's often the case that one is very close to the "nucleus" hires, including being a collaborator (I worked with Tzu-Mao Li early in his Ph.D. career at MIT, and worked very closely with Ren Ng [Berkeley] during his first two years as a Ph.D. student at Stanford, effectively in an advisorial capacity), or even being a formal mentor

(I was Manmohan Chandraker's postdoctoral advisor). I view this as a positive, ensuring the nucleus functions well together. Nevertheless, it is important that the additional faculty hires are further from your immediate orbit, and can bring new perspectives without any "in-breeding" risks. A last word is that it's not one size fits all. If you can't follow the recipe above, please don't hesitate to take advantage of opportunities when they do arise. The core faculty group at Berkeley (James O'Brien, Maneesh Agrawala and myself, along with older Profs. Barsky and Sequin) was not initially tightly integrated in terms of collaborations, but we still managed to become a powerhouse, and found collaborations among ourselves and elsewhere in the university once we were functioning as a group.

Obviously, one's role as a group leader doesn't stop when the group is hired. Earlier you were a king without a kingdom, now you actually have a group of junior (and in some cases senior) faculty. One needs to mentor the junior faculty, and this can be a significant effort. It is obviously easier when one is actively collaborating (one can give specific support on advising concerns if one knows the students involved, for example, and one can write joint grant proposals and give advice on that, one knows the work better to write award nominations and so on). But even with limited direct collaboration, you can be an effective mentor in helping the junior faculty member transition into faculty life, easing much of the initial frustrations, providing support, equipment and funding where available, and guiding them through the tenure review. Mentoring could be the subject of an entire chapter, but note that mentoring junior faculty into leaders can be as rewarding as mentoring of students or postdocs. While most mentoring discussion and effort is understandably focused on the pre-tenure years, faculty can benefit from advice and promotion throughout their career. If things go spectacularly well, perhaps you will transition from advising them on how to write their first grant proposal to thinking about the best way to prepare their Turing award nomination!

So far, we have given a recipe for success, but haven't discussed where to get the ingredients. Given the intense perceived competition for a faculty position, one may get the mistaken impression that all a hiring department needs to do is put out an ad, and faculty applicants will flock there. The truth of the matter is that it is extremely difficult to hire faculty, especially in a field like computer science, given burgeoning enrollments and faculty hiring from all departments, as well as attractive opportunities in industry. There's no panacea to this situation, but my general advice is to look beyond simply putting out an ad and seeing who applies. Be proactive and aggressive, seeking opportunity where others may not have looked. My model for this comes not from academia itself, but from the sports world. Imran Khan is now the prime minister of Pakistan, but when I was growing up, he was revered as one of cricket's greatest ever players and captains. More than his epic fast bowling or dynamic batting however, I was impressed by his team selection. The conventional way to select a national team in any sport is by looking at who is performing best in the domestic leagues. But Pakistani domestic cricket then as now, was moribund, and the statistics of players there did not reveal their technique and chances for success in the international arena. So, Imran Khan traveled around the country, trusting his own reading of a player's technique and Khan's ability to mentor players and find diamonds in the rough. From the outside, this looked like black magic, with teams regularly including unknown players essentially taken from "street cricket" (only a mild exaggeration), and consistently turned into international superstars.

The first faculty hire I made, Eitan Grinspun at Columbia, wasn't originally thinking of applying that year. But late in January, well past the time applications usually close, his postdoctoral mentor at NYU suggested that if he was thinking of staying in New York (which he wanted to do for personal reasons), he might as well apply now before the position disappeared. So, he sent me a non-committal e-mail, and to hear his telling of the story, I called back 30 seconds later... the rest is history. Ren Ng at Berkeley is of course an extraordinary individual who received the highest Ph.D. honor in computer science, the ACM Doctoral Dissertation award in 2006, but was applying for academia under unusual circumstances, seeking a re-entry into academia after nearly a decade at his startup Lytro, where he didn't publish at all. Manmohan Chandraker was happy in his appointment as the head of the vision group at the industrial lab at NEC Research. And Albert Chern was planning to apply only to math departments. None of these may even have applied to CSE academic positions in the US, leave alone the universities in question, without my urging, and they did not receive offers from other universities at the time they were hired. I view all of these, and others, as following the "Imran Khan" method, of finding talent wherever it may exist and be below the radar of other places, and then recruiting and mentoring them into academic superstars. That may well be my biggest legacy.

More broadly, my approach to having the biggest outreach is to contact any candidates I may know or have heard of, regardless of whether or not they may actually be on the market (what's the worst that could happen? They say no?) Then contact faculty elsewhere who may be of interest, perhaps hedging the e-mail as saying tell us who you think might be suitable, and of course if you yourself were interested, let's talk. For faculty we've considered hiring in the past but slipped through the cracks, I might be more aggressive saying we'd love to have interviewed you the first

time you applied, but we didn't have a slot in graphics then, we'd be happy to correct that wrong now if you wanted. I of course also reach out to important senior figures in the field to suggest candidates, either their own students or others. And keep track of opportunities that show up, by the candidates being proactive. Of course, if someone slips through all this, we ultimately make decisions based on the final list of applications, and as long as they have applied, we will consider them (if they apply elsewhere and show up on public interview pages, we have been known to contact them and ask them to apply as well). Eventually, one sometimes also needs to make a choice on whether to hire only "superstars" with proven track records (this applies to junior faculty, not just senior folks), or take your chances with bigger "diamonds in the rough." There is no one size fits all, one needs to be pragmatic based on the environment one is in. One cautionary tale comes from Columbia, where it was a widely known story that one research area focused only on "superstars" with a net result of hiring two faculty over a 17-year period, both of whom left in two years.

Finally, faculty hiring should not be viewed as a finite duration activity with an endpoint. That may well be the initial phase, such as establishing critical mass in the first 7 years at UCSD, with four junior faculty hires in CSE. But once that is established, one needs to continue hiring. Why? Fact of the matter is that faculty leave despite one's best efforts to create an ideal situation to support the entire group, or other priorities (family, startups, demotivation) take over in due course. As such, one should always be on the lookout for excellent faculty (although perhaps more opportunity-based than in the startup phase), and try to hire extraordinary faculty if possible, even if there is no immediate need. Again, a cautionary tale comes from both of my previous institutions. Columbia, given the nature of hiring in computer science, had slots strictly designated for research area. This is good, in that the area always has the slot, but bad in the sense that once we hired Eitan Grinspun, we could not grow the area further. So, once I left, graphics at Columbia was in a difficult place for a few years, then picked up steam again with the hiring of Changxi Zheng in 2012. But a few years later, when Eitan himself departed to Toronto, the situation fell back to being bleak, and that slot hasn't yet been replaced. A policy of continuous hiring in the intermediate years would have alleviated the backsliding when inevitable departures followed. Similarly, Berkeley grew into the leading group in the country during my tenure, and even hired Jonathan Ragan-Kelley (JRK) after my departure. But with losses of Maneesh, JRK and myself, graphics remains a shadow of its previous incarnation. Despite slots not being limited by research area, the competition from other areas has defeated the ability to rebuild. It's too early as of this writing to say how the future will play out at UCSD, and if we will fall into a similar trap, but it will not be for the lack of my trying.

Returning to the analogy at the start of this chapter, you now have your ships (or at least your ship captains), for your expeditionary force in your voyage of exploration into the unknown. But once you set sail into the tumultuous ocean, how do you ensure that your armada (or faculty group) functions in a cohesive and effective way? The next chapters discuss a number of issues involved in building the lab or center, and keeping the faculty and student group engaged and active.

3. LABS AND CENTERS

Shortly after I arrived in New York City in Fall 2002 to take my first faculty position at Columbia University, a group of seven professors on the seventh floor of the CEPSR building met to discuss the formation and branding of a new center in computer vision and graphics. A flippant suggestion for the name did involve calling the center Viagra, for VIsion And Graphics. But ultimately, we settled on the rather unimaginative but unambiguous “Columbia Vision and Graphics Center” or CVGC, appointing Profs. Shree Nayar and Steve Feiner as co-directors (titles they still hold). There immediately followed a great burst of enthusiasm to develop some logos and branding, a CVGC lecture series, signage at the entrance to the floor, and attempts to build a sense of community, a brand, a funding base and a model for student admissions and mentoring. Some of these activities were more successful than others, and over time some of the initial enthusiasm faded. Nevertheless, CVGC lives on as an umbrella for a strong group of faculty in vision, graphics and related areas at Columbia, even adding more faculty over the years. I have been fortunate to subsequently participate in the creation of the Visual Computing Lab (VCL) at UC Berkeley, and the UC San Diego Center for Visual Computing (VisComp) at UCSD. All of these centers have played an important and ongoing role for their students and faculty, and indeed it is hard today to think of any of these universities without them. Why are centers and labs important, how are they to be created, and what are best practices in their management? This chapter addresses these topics, discusses how a group leader can build and facilitate centers, the role of faculty and students, and how university administration can assist with their creation and recognition.

At the outset, I would like to note that labs and centers are amongst the most over-used terms in academia, with no clear definition. A lab could stand for anything from a single faculty member’s own group and research laboratory to an entire department (for example, MIT’s CSAIL or Computer Science and AI Lab). Similarly, centers can mean anything from a grant-funded research project to a building-sized group of faculty. In this chapter, we view labs and centers as referring to faculty at a university in a specific area (such as vision and graphics, although the group considered in a center is usually more inclusive and involves more than the immediate core faculty). Within our definition, a center will continue indefinitely. Of course, other models exist... for instance the project-based centers within the systems group at Berkeley, or larger multi-faculty centers supported by the NSF or industry, but these are usually not branded within a particular research area and institution, and cease to exist once the particular funding source is exhausted. Similarly, for a smaller group of two or three faculty, simply calling it something like the “graphics group” may be adequate, and there may be no need to define a broader entity such as a lab or center.

I view research-area based centers as having two main purposes. First is simply one of branding and citizenship. It is an opportunity to advertise to the world that your university has a large group of faculty working in a particular research field, including adjacent topics. This can provide a stronger sense of coherence than simply viewing the research as a union of the output of individual faculty member’s labs. Faculty can view themselves as “citizens” of the center, with which they may have a stronger affinity than the department or the university. This is especially the case for graduate students, both when deciding which university to join for their degrees, and in identifying their affiliation within the university. While undergraduates usually view their primary loyalty as to the university and secondarily to the department of their major, graduate students often aren’t really looking at a university as a whole (indeed, witness the power of departmental rankings, but there is nothing really like the national ranking of undergraduate universities at the graduate level). But in today’s world, incoming students are already specialized with prior research experience in the areas they want to study. As such, a department is also too broad, and they are really applying to a specific research area or center. At the other extreme, while some students may be set on working for a specific professor, that seems too narrow in their sense of loyalty within the university. Therefore, the group, lab or center is often the right granularity in terms of defining where students’ “citizenship” lies, and the overall strength students consider when choosing a university to pursue their higher studies at. Graduates of our group at UCSD often refer to themselves as graduates of the Center for Visual Computing (VisComp), even though formally VisComp confers no degrees (those are awarded by the department and university).

Second, the center forms an administrative umbrella under which a whole set of activities runs smoothly. For example, we can invite external speakers and have an internal seminar series, which is a lot more coherent than having ten different small seminar series for different faculty, and having to participate in ten different external visits. We have a natural affinity group to discuss course scheduling and our curricular programs. We have a natural group to organize a visit day for incoming Ph.D. students in the research field. If you think carefully about the complexities of academic life, there are a myriad such situations. The center can also facilitate funding, especially from industrial sponsors (we regularly have an annual retreat with industrial sponsors). A well-run center will be well-recognized externally, not

just within the university. Just as your own students view their citizenship as lying primarily with the center, external researchers will mark students and faculty as members of your center.

We've seen the benefits of having a center. So, how do you create one? At one level, it's extremely simple. Once you've hired your faculty group, as discussed in the previous chapter, all you need to do is get the relevant faculty into a room (or in the modern age on a zoom call) and agree to do it. It is beneficial to have a greater critical mass than just the core group of three or four faculty you are working most closely with, and this also enables the center to be created even before or while you are in the midst of faculty hiring, serving as a powerful attractor for new faculty and students. We had eight faculty in the CVGC at Columbia after Eitan Grinspun's hiring, five in the VCL at Berkeley, and a dozen in UCSD's VisComp across four different departments. Of course, simply naming a center will aptly get you a "center in name only." Not that this is necessarily a bad thing; having something in name only is better than having nothing at all, as it enables one to get some of the branding and administrative benefits explained earlier with very little overhead. In what follows, I describe some of the aspects required in going beyond that to obtain a more functional group and lab or center (for the most part, I use these terms interchangeably).

Any center needs some form of external signage. In its simplest form, this involves a well-maintained website (at Columbia, we paid a significant amount for a design firm, which ultimately botched the design and did not work well with academics, not a path I would recommend for the future). It is also a good idea to have some physical signage at the entrance to the center space or group of contiguous faculty. At UCSD, we spent several years and over a hundred thousand dollars renovating a physical laboratory for the center for visual computing, which served as a frontispiece for our group, as well as providing a spacious location for students, talks and research projects. At Stanford, Columbia and Berkeley, we also printed large frames of images produced by previous research projects that were displayed prominently (at UCSD, we upgraded to using large display monitors with loops over images and videos). Other groups have been known to print commemorative t-shirts and display other swag. While I could go into entire chapters on each of these topics, many of the details and best practices are quite specific to a particular institution. Besides the physical infrastructure, one needs to provide for the intellectual life of the center, including both internal and external seminar series, and perhaps monthly or quarterly dinners with members of the center (Stanford has even been known to have annual ski trips and lab excursions). An annual retreat for the center is a good idea from both a social and a research perspective (we have ours at the La Jolla Shores hotel, with 60+ participants). Needless to say, each of these activities involves considerable effort and organization to set up and maintain, most of whose burden will fall on the group leader or center director. But the rewards for doing this can be significant, and important to maintain a visible internal and external presence that students and faculty will be proud to be part of.

The center can also be a vehicle to acquire and maintain relevant research equipment, depending on one's research area, as well as other general purpose supplies and resources such as printers, 3D printers, cameras, and even a compute cluster for intensive work, all of which we maintain and provide at UCSD. This then leads to the natural question of funding and how these "common good" resources are to be paid for. It would certainly help to create a more cohesive group of faculty and research program if one could raise external funding. In some cases, this may be possible through large-scale ("center-scale") proposals to federal funding agencies. In other cases, one can approach industry to have a robust industrially-funded center, either through support of a particular research agenda, and/or in terms of a membership fee model, where a select group of companies pays an annual membership fee (\$50K-\$200K), which is then used for center "common good" activities, administrative support, and the faculty, with excess distributed either as fellowships to students or to specific research projects based on a lightweight internal proposal review of one to two page submissions. Perhaps the most successful example of this last model, which has formed my own inspiration, is the Berkeley Vision and Learning Center started by Trevor Darrell (now Berkeley Artificial Intelligence Research or BAIR), which grew to a consortium of 20+ companies. You don't need to be as successful as Trevor to maintain a functioning center, but having some unrestricted membership fee-based funding can be extremely helpful in lubricating the functioning of a center and paying for a number of the shared resources discussed above. Note that shared resources include activities such as reimbursing speakers' travel for seminar series, paying for food for internal seminars, the annual retreats, and social dinners or lunches. Without a funding source, these activities would either need to be stopped, or one would have to request contributions from individual faculty to be pooled to pay for them. While that approach can work, it can lead to individual faculty claiming poverty and opting out, or feeling resentful, and in any event lead to friction within the group. As a group leader, my advice is to either develop an external funding model or if you are otherwise well funded, consider just paying for at least the smaller activities yourself. It's a simple magnanimous gesture that is well rewarded in terms of bonhomie within the group.

It is important to recognize that the university needs to support these endeavors, after all they control the slots for faculty hiring in the previous chapter, and a group leader will often require some measure of university support and resources to build a functioning center. Yet, in one way, everything I've written so far is a dream. There is no formal position of "group leader" or "center director" at most US universities. The only formal organization is into departments (such as computer science), and schools (engineering), and the only formal administrative titles are department chair, dean and provost/president. The administrative structure for faculty is also flat. They are independent, without any notion of a "boss" and to the extent they report to anyone, it is only to the department chair and dean. They have no formal requirement to listen to anything a group leader or center directory says, and you have no formal power over any of the faculty in the center. As such, any progress you can make must come from the shared vision and consent of other faculty in the center, and it is important to hire faculty who share this vision, can work well together, and recognize you as a highly regarded mentor and leader. Ultimately, many groups and centers begin a whole slew of activities with great fanfare, but it is hard to sustain the level of enthusiasm year over year, and the cohesion can die out simply from apathy and the general busy life of faculty, rather than any explicit decision. In this context, having an externally funded center as discussed in the previous paragraph can help, as faculty are always more likely to continue to participate when there is money involved.

University administrators like the chair, dean and provost need to understand that while being a group leader is not a formal title, it does carry real opportunities and responsibilities, and should be supported or regularized into an actual administrative position if possible. And the university ultimately benefits significantly when a highly visible center is created. Indeed, the foresight of former Columbia provost Jonathan Cole and dean Zvi Galil was a major catalyst for the creation of CVGC, and the vision of current UCSD Dean Al Pisano was the main impetus and support for the creation of VisComp. But even without the "push" from enlightened administrators, as a group leader, you can "pull" in the administration to get the resources and recognition you need. It is beneficial for one to be conferred a formal title as "center director" or equivalent to provide some recognition of the significant administrative role, even better if this comes with some real benefits, such as teaching relief or an administrative stipend. Beyond this, negotiations with administration often concern faculty slots, space, equipment, administrative support and funding. Truth be told, these negotiations are easiest to have during recruiting and retention of the faculty group leader, but it is important for administration to be receptive even during normal circumstances. And it is certainly worth considering a potential and proven track record for building research groups within the hiring process for senior faculty, and success and leadership in building a group or center within merit reviews, even if it doesn't easily lie within the conventional pillars of success in research, teaching and committee service.

In many ways, a multi-faculty research group serves as one's extended professional "family." Describing it as a large Indian "joint family" is perhaps more apt, being an agglomeration of several individual faculty research groups. A well-functioning group or center can be extremely powerful as described above, amplifying the research output, providing a sense of community, and even serving as a powerful retention tool; after all people are loath to leave their supportive family for foreign pastures. But we must also deal with the truth that some families are more functional than others, and dysfunction, whether in joint families or research groups and centers, is common. Faculty within the center may be at loggerheads with each other, simply apathetic, disagree on everything, or even no longer be on talking terms. Under these circumstances, participation in a center may be impractical, and indeed one may view the larger research group as a millstone around one's neck, rather than a positive amplification of one's impact. As such, while I hope you as a group leader or a faculty member are never in a dysfunctional lab or department, this chapter does need to deal with those issues.

As a group leader, you seek to hire faculty who are not just research stars but will work well together, often described in part as a "no jerks" rule. For example, one might explicitly choose not to pursue senior hires who have reputations for being difficult to work with, independent of their research caliber. One tries to ensure that junior faculty exhibit no red flags in their behavior at faculty interviews that may mark them as an undesirable colleague, and it's of course easier if one already has a working relationship with them. This can cut both ways; you are more likely to hire colleagues who want to work with, and get along well, with you. When maintaining the group, pay attention to ensure things run smoothly. If conflict breaks out between other colleagues in your lab or center, you may well be called on, and should be proactive as a mediator to smooth over this conflict and find a good solution for all. It helps if you have the support of your colleagues and can be trusted as a fair mediator who always has the best interests of the group and its faculty in mind. In this context, it's usually helpful to have a single well-respected group leader, than to have leadership effectively by a small council of faculty, who may each have very different views of how things should run. Finally, it is advisable to be generous as a group leader or as an individual faculty member. In any center or lab,

there are likely to be those who contribute more than others to shared resources versus the “free-loaders.” For the harmony of the group, it is best to simply pay for common items such as equipment, retreats, lunches, administrative staff, printers from your own funds or those of your center, rather than trying to ensure faculty pay in proportion to their usage. The amount of goodwill you generate, and the level of friction you avoid is well worth whatever price you think you may bear in terms of fairness and free-loading.

But all this planning can only go so far, and you may be unable to avoid existing tenured faculty in your research area with whom you have a difficult relationship, or you may join an existing group, either as a leader or as an individual faculty member, which already has strong frictions. What can you do then? One option is for “divorce” as in family court. In the extreme case, this may involve changing institutions or departments. But there are simpler alternatives. By being somewhat creative in how you define your research field or center, it may be possible to either join another more functional research group or area, or to define the research direction of the lab or center in such a way that the “troublemakers” are excluded, and you can create a strong coalition of engaged faculty who will actually work together. On other occasions, you may feel you are better served by not joining a center or lab, but instead working as an individual faculty. However, this may also not be practical if you are strongly labeled as being within the research area, and for administrative reasons. Imagine if everyone started doing this. A department chair is not going to be thrilled about dealing with 50 individual faculty rather than 10 groups, leave alone the negative impact it would have on department morale. It’s similar to “divorcing” a department; there is no procedure to be an unaffiliated faculty independent of any department, and indeed the president and provost wouldn’t be thrilled about individually managing 2000 faculty rather than a few dozen department chairs. Ultimately, your only practical option may be “benign neglect” [I acknowledge Jitendra Malik for coining this term in a different context]. In this mode, you passively stop participating (“neglect”) in the center or lab activities, or participate minimally in only those activities of direct benefit to you, while still maintaining civil relations (“benign”) with the group leader and others. Ultimately, you can continue collaborating with only those faculty with whom you work well in research terms, and you can leave the center to its own devices, without necessarily damaging any relationships explicitly. Of course, I hope you never run into these issues, and the early part of the chapter serves as a guide to establishing a fully functional research group and center.

All of the faculty hiring and center-building you can do will have little impact without supporting the core mission of the university, achieving excellence in research and teaching. Ultimately, these are responsibilities of the individual faculty members involved, there are many guides for faculty on the market, and I can’t pretend to tell you what research to do. But there are things a center can establish to facilitate and support the research and teaching of its faculty, and indeed things a faculty group should do to have a strong research and teaching program, ably supported by a strong group of graduate students. This will be the subject of the next chapter.

4. RESEARCH, TEACHING AND STUDENTS

When I was in graduate school and striving to receive some external advice about the graduate experience, I read Peter Feibelman's excellent book, "A PhD is Not Enough! A Guide to Survival in Science." In his chapter on establishing a research program, Feibelman makes clear that while he provides much excellent advice, "I am not going to tell you how to win a Nobel Prize... I can't even tell myself how to do so." In the same vein, this chapter is not really about how and what to research, but rather a discussion of how a center or multi-faculty research group can support the careers of all its members, and build a stronger research and teaching program.

A strong group of inter-connected faculty can facilitate many collaborations amongst them, and indeed over time one would expect to see collaborative papers between any two professors. This has been one of my most satisfying experiences at Columbia, Berkeley and UCSD. (I should note for the record that this is not the only model, and very strong research groups can be achieved by individual faculty demonstrating excellence in their own sub-fields while working cooperatively together, but not necessarily collaborating on papers). Having faculty with similar interests also enables joint grant proposals, and a shared mission for equipment, space and graduate recruiting. The last is the issue I want to focus most on, since the quality of graduate students ultimately has a substantial impact on one's career.

When I got to Columbia, I think it is safe to say that in my first year there (2002-03), we had the worst possible Ph.D. admissions system in the country. Applications were still being received on paper folders (or printed that way from graduate division). They would come to us on the admissions committee in small batches of 10 or 15, with most applications actually not received in the department until after the national candidate reply date of April 15. Needless to say, most students you would want to admit had already accepted offers elsewhere, or were on the verge of doing so, before we even looked at their applications. There was certainly no concept of a "visit day" to let students visit the department, while we tried to recruit them. Faculty were given complete control of who they wanted to admit individually, but with a catch. Any offer to a student required full funding from the faculty member's grants for at least the first year, and there was no yield insurance... you might expect only one in three students you admit to actually accept the offer, but you had to have the funding for all three (and if nobody turned up, your projects would simply not get done). Realistically as a student, you weren't going to get admitted to Columbia Computer Science unless a faculty member knew of your application outside of the formal admissions process, and unless you were willing to already commit informally to joining prior to receiving a formal offer of admission. Incidentally, this experience still leads me to give contrarian advice to students applying to graduate school. Many professors advise students not to contact them before admissions, as there is little that can be done until they see the full application and let the formal procedure take its course. On the other hand, I always advise students to email me or other faculty of interest prior to applying, as one should make no assumption that submitting one's application by the formal procedure actually brings it to the attention of faculty of interest, or for that matter that anyone will read it at all.

Somewhat paradoxically, reforming a system is often easier when it is completely broken, rather than when you simply want incremental change. After the experience above, nobody on the Columbia faculty would support continuing the current system, and over the next few years, we created among the best run admissions processes. The first step was to implement a modern electronic system for CS Ph.D. applications, courtesy of then department chair Henning Schulzrinne. Thereafter my colleague Eitan Grinspun, my first faculty hire in 2004, organized a very well-run department-wide visit day, which substantially boosted recruitment yields. Admissions applications were received early, and we made decisions early, competing on even terms with other top schools (and aided substantially by the presence of a strong collaborative group). The one remaining issue was the question of funding first-year students. The challenge of "yield-protection" remained, viz. you want to make multiple offers for a given slot given the probability students accept, but you don't have funding to guarantee a larger than expected yield, especially if a faculty member needs to fully pay for a student. Moreover, many students haven't matched with a single faculty member prior to arriving in the department. Indeed their "citizenship" is usually within the larger center, as discussed earlier. Giving them the opportunity to be co-advised, do rotations, or pick an advisor later is advantageous in many cases.

Ultimately, the system we came up with was one of fractional or shared funding. On any student, if multiple faculty were interested, they would split the costs, at least for the first year. This allowed us to make more offers (since one's risk is reduced in terms of the funds committed to any student), it gave the students more choice making them easier to recruit, and it got more faculty into play and committed to each student recruitment. And if a student clearly matches to a single professor only, or any individual faculty member doesn't want to participate, the old system can still be used, so there is little disadvantage. Our method had advantages well beyond the initial admission. Most students

initially co-admitted and co-advised by multiple faculty continued that arrangement through their Ph.Ds facilitating collaborations between faculty, a more collaborative lab culture, and a reduced funding requirement for all faculty. Ultimately if a large number of one's students are co-advised and split-funded, the amount of funding one needs to raise is less. And even when writing joint grants, funding agencies will usually scale the funding for the number of faculty. Throughout my career, I have routinely paid less than half of the cost for my students (helped by fellowships and required teaching assistantships in addition to co-advising). Needless to say, a more collaborative environment also enables new problems to be solved, beyond the expertise of any individual faculty member.

The same story has repeated at other universities. Berkeley often had a great deal of difficulty in terms of Ph.D. admission because of the funding and admission model, but Trevor Darrell ultimately largely resolved that problem by setting up the Berkeley Vision and Learning Center for the initial purpose of providing fellowships to all first year admitted Ph.D. students in the area. UCSD CSE is well-enough endowed to support most of the first years with fellowships, with additional yield-protection insurance, but even so our discussions often center around making the funding and admissions model even better. More collaboration and co-advised students can even make one a more effective mentor as a lab director. If one is actually co-advising students and directly observing a junior faculty member (and their students), one can give more effective and precise guidance on advising, problems with students, and other issues. In general, any strong group will want a strong Ph.D. admissions system with associated funding, a strong mechanism to recruit those students, and ideally mentoring within a collaborative broader research group. Indeed, students will be attracted to a friendly collaborative group of faculty, often preferring that to individual faculty members elsewhere who may otherwise have attractive research programs or be in higher-ranked departments.

Like in the last chapter, I must acknowledge that not all groups may function in a harmonious fashion. The advantage of the approach I've outlined is that it is voluntary; you are not required to collaborate with any faculty member you do not wish to, or admit jointly advised students. You are welcome to strike out on your own if needed. Nevertheless, friction may arise, especially if some faculty are perceived as "free riding" on the system, not contributing their fair share for first-year students who ultimately end up working with them. Once again, I believe as a group leader, one should not try to enforce absolute fairness. Be generous, allow some free riding if needed; ultimately having more resources and more strong Ph.D. students makes everyone's life happier, and enables great research and great students, for which your lab or center will ultimately be known.

Beyond research, there are numerous other functions a group must facilitate. A key aspect of this is the curriculum and courses. The members of a center should get together and plan a coherent program of courses that its faculty will offer, with an equitable distribution of teaching load. It is best if the faculty do this internally, and then get the department to rubber-stamp approval, rather than requiring the department to intercede and enforce a plan. There should be some discussion of redundancies if an individual faculty member is on leave, sabbatical etc. and perhaps a group of occasional lecturers who can pick up some slack if that's the case. Moreover, core introductory courses should ideally have more than one faculty member who can teach them regularly to help out in such cases, and to provide the individual faculty with more variety. Ideally, faculty may even be able to combine forces to create a joint set of lecture notes, slides, assignments and MOOCs on core topics to enable continuity in teaching internally, and to promote the educational product to the world via textbooks and MOOCs. At UCSD, we were able to launch the first multi-course professional certificate on virtual reality on edX by Prof. Schulze and myself. It is increasingly the case that research groups will be known by not just their research leadership, but also by their educational leadership. I am perhaps proudest in terms of awards of my two edX prize certificates for exceptional contributions in online teaching and learning for the introductory computer graphics course. At Stanford, the PBRT book that came out of the physically based rendering course became the first (and to date only) book to receive a technical academy award. And courses serve as the testbed for introducing new ideas especially at the graduate level, and identifying promising research students to then be mentored into successful graduate students.

Speaking of students, I encourage all faculty and groups to be open to possibilities beyond the regular Ph.D. admissions process. Many universities have master's programs, which can be great sources of students, who can usually move with few fetters into the Ph.D. program upon identifying a willing mentor. Indeed, teaching graduate courses and exciting those master's students is the best synergy of research and teaching. Some of my best students in the Columbia years came through this mechanism. It is also fulfilling to include senior and even early-stage undergraduates in research, and in some cases, those too can become Ph.D. students (although for undergraduates, most advice holds they are best off moving elsewhere for graduate school). Ultimately at a university, all of the research and teaching must be coupled with producing our greatest product, which is the people that graduate from

our programs. My former department chair at UCSD, Rajesh Gupta once said “Our most important products are our people.” My undergraduate advisor Al Barr when giving me sage counsel upon starting my first faculty position commented that if one does one’s job well, “you will be known, not for the research you do, but for the students you graduate.” As a lab director, you have the fortune that this applies not just to students but also to the junior faculty you will hire; seeing them blossom into leaders will be amongst your biggest rewards.

5. EPILOGUE

In March 2020, the coronavirus pandemic abruptly shut down departments and forced universities to switch to remote operation. Universities were suddenly faced with dire economic and health consequences. Faculty at UCSD, being required to transition on a dime to remote teaching, were also now required to designate surrogates to maintain “teaching continuity”, the unspoken assumption being that any instructor may drop dead at any time. At the same time, faculty hiring came to an abrupt halt at many departments across the country, research in many experimental labs was brought to a complete standstill, and universities faced unprecedented financial losses with the usual concerns of a stock market recession impacting endowments and loss of state support coupled this time with the possibility of massive losses in tuition, and certainly in housing and dining revenue.

The pandemic no doubt led to close friends, family and acquaintances suffering and dying for nearly all of us, and continues to have substantial economic and health impact. Nevertheless, the worst fears in those early days of March 2020 did not come to pass. Within the context of leading a research group, which is the focus of this article, my goal was to find the opportunities in this crisis and move forward despite the headwinds. Indeed, we achieved many of our strongest milestones over the 18 months of remote operation. As one of the few universities that continued unabated with hiring plans through spring 2020, we were richly rewarded when Prof. Tzu-Mao Li accepted the offer to join our group (starting in Fall 2021), completing the initial batch of four Computer Science hires, and in some ways marking an end to at least the first stanza of my adventure in building the visual computing group at UCSD.

As remote operation continued into 2020-2021, the graphics and vision group at UCSD achieved new levels of productivity even in the midst of a global pandemic, with 25 papers each at the leading computer vision conferences in 2021, CVPR and ICCV, and 10 papers at the leading computer graphics venue, SIGGRAPH 2021. In terms of my own research too, this period marked a watershed. 2021 was the most productive year of my career to date, and my 2020 paper (with Ben, Pratul, Matt, Ren, Jon) on Neural Radiance Fields or NeRFs was already close to picking up its 500th citation. In fact, Frank Dellaert at Georgia Tech created a superb website to keep track of NeRF extensions, writing “Besides the COVID-19 pandemic and political upheaval in the US, 2020 was also the year in which neural volume rendering exploded onto the scene, triggered by the impressive NeRF paper...” The metrics-based website csrankings.org ranked visual computing (vision and graphics) at UCSD #1 for the most recent three-year period 2019-21, as I write this at the end of October 2021.

As opposed to some of the events recounted at the start of this article, these developments proceeded with no public fanfare, derogatory comments, discussion or recognition. As noted in earlier chapters, being a group leader is not an official title, nor are there any specific awards for those who are successful in this endeavor; it’s not even an official criterion for merit advancements within the university. This is not to say these achievements went unnoticed in the field. Kavita Bala, dean of computing and information sciences at Cornell did note to me “Congratulations on getting Tzu-Mao at UCSD. Wonderful group you [got] going! ... Hope we can catch up some time, though with everything going virtual and up in flames, who knows when that will be.” Similar sentiments have been expressed by a number of recent virtual and physical visitors to the department. In terms of research-related awards, I was happy to see the recognition deservedly go this time to the students who did the bulk of the work. In the past four years, four of my Ph.D. students won at least department research or thesis awards, and recognition within the field or the university, including the SIGGRAPH Outstanding Dissertation Award, and the UCSD Chancellor’s Dissertation Medal, while other collaborators received a couple of other SIGGRAPH Dissertation and New Researcher Awards.

I felt truly fortunate to have participated in, and indeed led, this remarkable journey, both at UCSD over the past 7 years, and over my career to date, while at the same time looking forward eagerly to the second half of my career. Indeed, to quote Steve Jobs, “The Journey is the Reward.”