MIT EECS
ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

CSAIL

Nov. 29, 2018

Bill Freeman, Antonio Torralba, Phillip Isola

6.819 / 6.869 Computer Vision

How to do research

Slow down to speed-up

The jump from problem sets to research can be hard. We sometimes see students who ace their classes struggle with their research. In little bites, here is what I think is important for succeeding in research as a graduate student.

The first advice can go on a bumper sticker: "Slow down to speed up". In classes, the world is rigged. There's a simple correct answer and the problem is structured to let you come to that answer. You get feedback with the correct answer within a day after you submit anything. Research is different. No one tells you the right answer, we don't know if there is a right answer. **We don't know if something doesn't work because there's a silly mistake in the program or because a broad set of assumptions is flawed.** How do you deal with that? **Take things slowly.**

Bring the problem back to something where you know what the answer should be. Verify your assumptions. Try something really simple. Understand the thing, whatever it is–the program, the algorithm, or the proof. **As you do experiments, only change one thing at a time, so you know what the outcome of the experiment means.** It may feel like you're going slowly, but you'll be making much more progress than if you flail around, trying different things, but not understanding what's going on.

“It doesn’t work”

- Please don’t tell me “it doesn’t work”. Of course it doesn’t work. If there’s a single mistake in the chain, the whole thing won’t work, and how could you possibly go through all those steps without making a mistake somewhere? What I want to hear instead is something like, “**I’ve narrowed down the problem to step B. Until step A, you can see that it works, because you put in X and you get Y out, as we expect. You can see how it fails here at B. I’ve ruled out W and Z as the cause.**”

Working hard

“**This sounds like hard work.**” Yes. It’s no longer about being smart. By now, everyone around you is smart. In graduate school, it’s the *hard workers* who pull ahead. This happens in sports, too. You always read stories about how hard the great players work, being the first ones out to practice, the last ones to leave, etc.

Love what you do

“How do I get myself to work hard enough to do research well?” It all plays out if you love what you’re doing. You become good at it because you spend time at it and you do that because you enjoy it. So pick something to work on that you can love. If you’re not the type who falls in love with a problem, then just know that working hard is what you have to do to succeed at research.

Steering

I have to note that the above isn't completely true. Beyond working hard, there's also *steering*. We're like boats. We need motors–that's the working hard part. **But we also need a rudder for steering–that's stepping back periodically to make sure we're working on the right thing.** On the topic of steering, I find time management books to be very helpful. They teach you how to spend your time solving the right problems.

Toy examples are helpful tools

There's a concept I want a simple phrase for, and maybe you can help me think up a good name. It's the simplest toy model that captures the main idea. TSTMTCTMI ? Anyway, simple toy models always help me. With a good one, you can build up intuition about what matters, which is a big advantage in research.

Here's an example. The color constancy problem is to estimate surface reflectance colors when we only get to observe the wavelength-by-wavelength product of the each surface reflectance spectrum and the unknown illuminant spectrum. A toy model for that problem is to try to estimate the scalars a and b from only observing their product, y = ab. There's a surprising richness even to this simple problem, and thinking about it allows you to think through loss functions and other aspects of Bayesian decision theory. I co-authored a paper that discusses y = ab for much of the manuscript. Another toy model: as a proxy for complicated shaded surfaces, a single bump. You get the idea. Having the intuitions from working

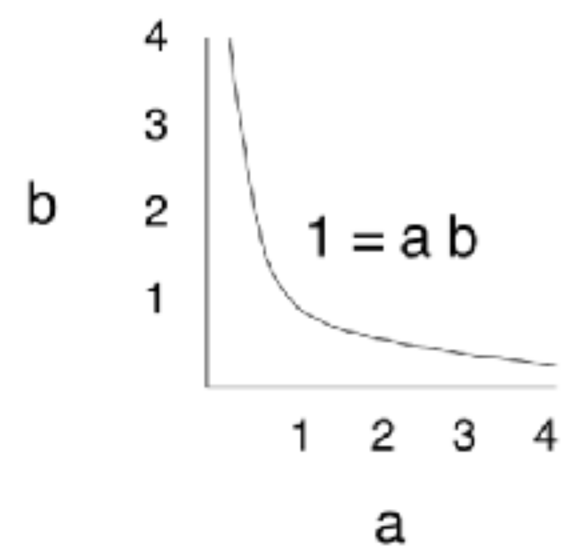

Figure 2: A toy model for the color constancy problem: y = ab

The parable of two students

A parable, as told by my friend Yair Weiss: **There is a weak and a strong graduate student. They are both asked by their advisor to try a particular approach to solving a research problem.**

The **weak student** does exactly what the advisor has asked. But the advisor's solution fails, and the student reports that failure.

The **strong student** starts doing what the advisor has asked, sees that it doesn't work, looks around within some epsilon ball of the original proposal to find what does work, and reports that solution.

Ways to think about other researchers

- **Sometimes it’s useful to think that everyone else is an idiot.** This lets you do things that no one else is doing. It’s best not to be too vocal about that. You can say something like “Oh, I just thought I’d try out this direction”.

- **On the other hand, it’s also useful to remember that many smart people have worked on this and related problems and written their thoughts and results down in papers**. Don’t be caught flat-footed with a large body of closely related literature that you aren’t familiar with.

Your brand

- Here’s how a business school might talk about your research. You have a brand: you. There are many impressions you want to build up about your brand: that person always does great work, they have good ideas, they give great talks, they write wonderful software. Promote your brand by being what you want to be seen as. Build up a great reputation for yourself.

- Cultivate your strengths and play to those strengths. Some possible strengths: being broad; creative; a great implementer; great at doing theory.

Nurture your research brand.

Reporting progress to your advisor

- Please don’t report to me, “This instance doesn’t work”. Why doesn’t it work? Why should it work? Is there a simpler case we can make it work? Do you think it’s a general issue that affects all problems of this category? Can you think of what’s not working? Can you contort things to make an example that does work? At least, can you make it fail worse, so we understand some aspects of the system?

- I love to hear about progress when I meet with students, but note that I have a very general notion of progress. **Progress can include: “I’ve shown why this doesn’t work”, “I’ve simplified the task to get it to start working.”, or “I spent the whole time reading because I know I have to understand this before I can make any progress.”**

- Please don’t hide from me. Let’s talk. I like it when you track me down and insist that we talk, for example, if I’ve been traveling.

Final note

One final note about doing research: I hope you love it. I certainly do. The research community is a community of people who are passionate about what they do, and we welcome you to it!

Polaroid 1981-1987

China 1987-1988

Time

PhD 1988-1992

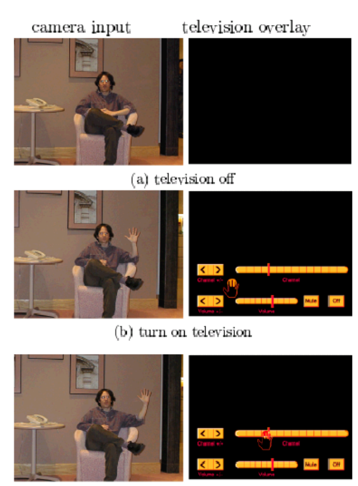

MERL 1992-2001

MIT 2001-present

Elements of a successful graduate career

I crowd-sourced the rest of this talk. I sent this e-mail to the MIT Computer Science faculty and other CSAIL researchers:

> Could each of you please send me what you think is the most important quality for success in graduate school?

Following are the answers I got back.

Desired qualities

The most important qualities: curiosity and creativity

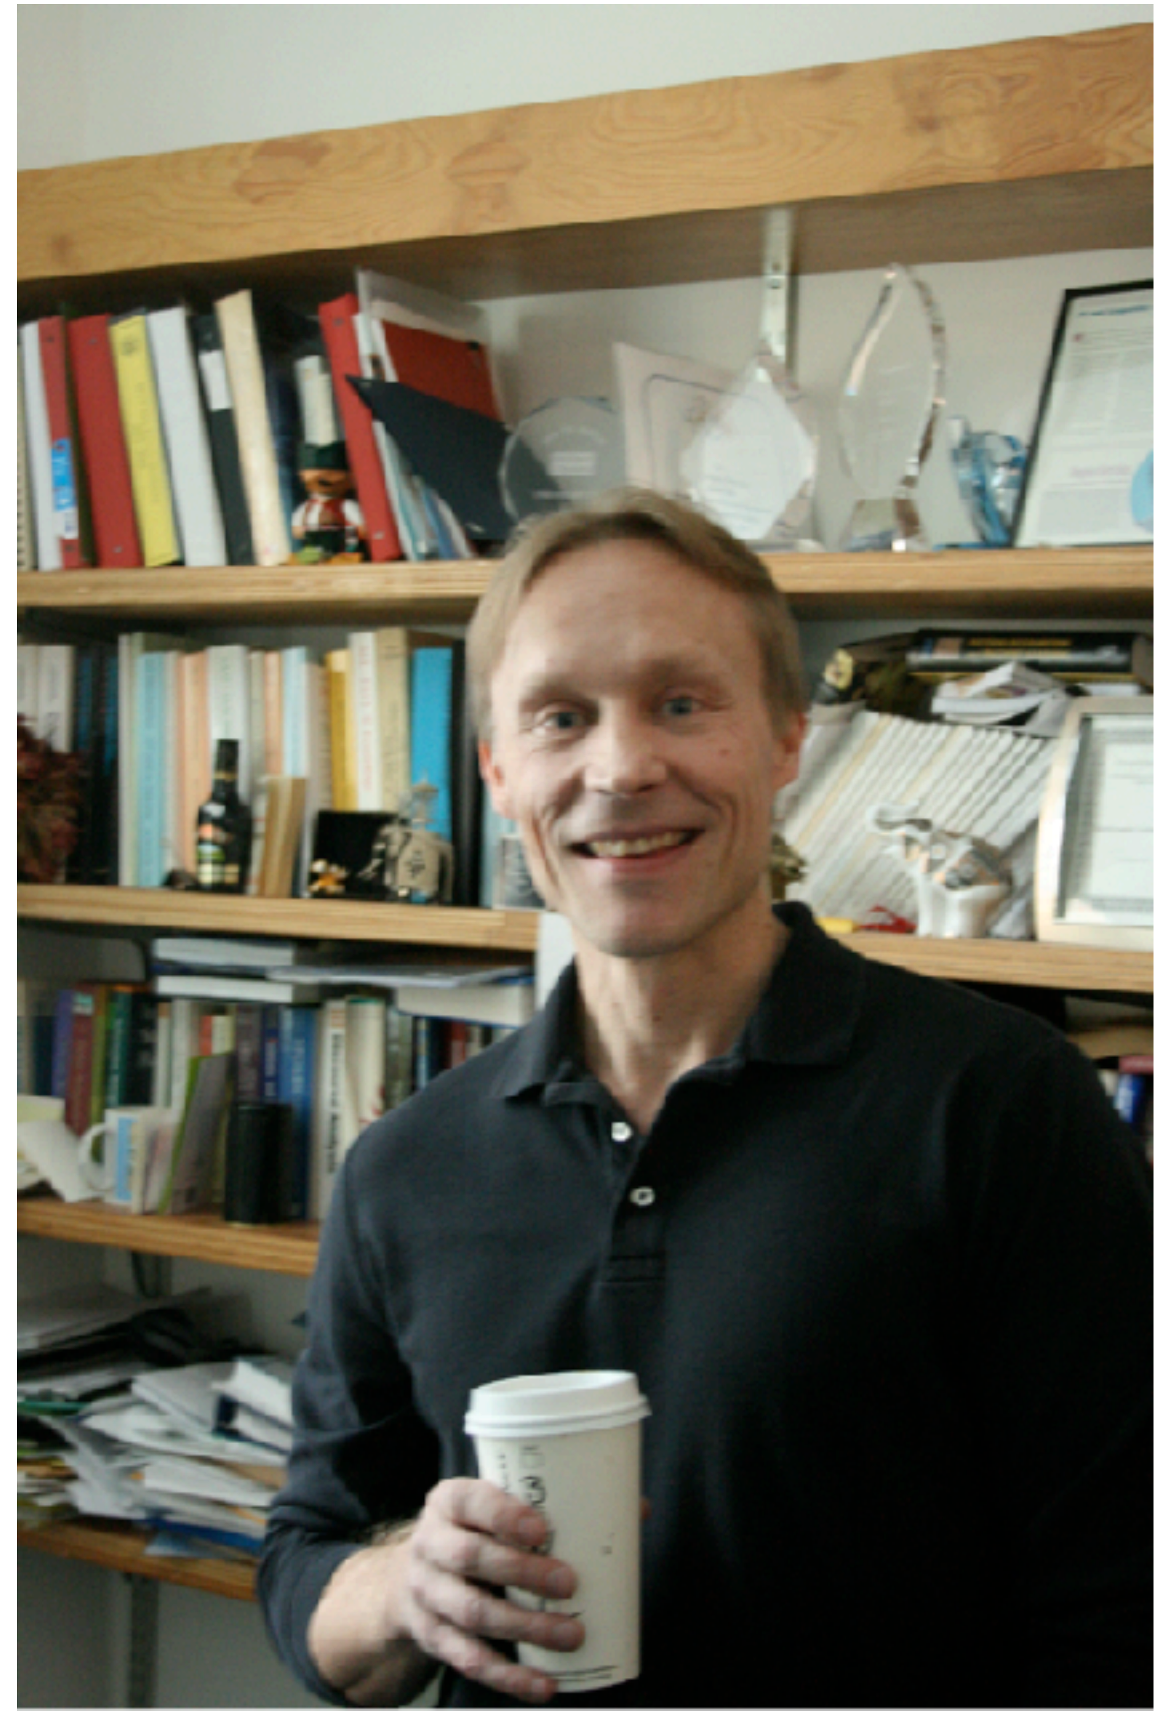

Tommi Jaakkola,
Machine Learning

Here are my desired qualities:

Determination, curiosity and flexibility, willingness to work hard,

try not to be ruled by fear...

Shafi Goldwasser
Cryptography

You want to be a marathon runner, not a sprinter.

(Not necessarily just the good, quick answers, as in a class. Want to be able to come up with the more thoughtful answers, too).

Brian Williams,
Robotics

David Karger
Algorithms, HCI

Purposefulness.

You are in grad school for a purpose (whatever your purpose is) and it is up to you see that purpose accomplished. Now, when nobody's telling you what to do, you have to tell yourself.

Dave Gifford,
Bioinformatics

Total fascination with an area of science coupled with drive and imagination!

Jack Dennis,
Computer Architecture

My first thoughts are: persistence, courage, flexibility

My choice for most important: persistence.

Pete Szolovits,
Healthcare

Enthusiasm, curiosity, lots of energy, and scholarship. Being smart also helps, but admissions has already taken care of that.

Advice

Enjoy the freedom you have in graduate school to be creative

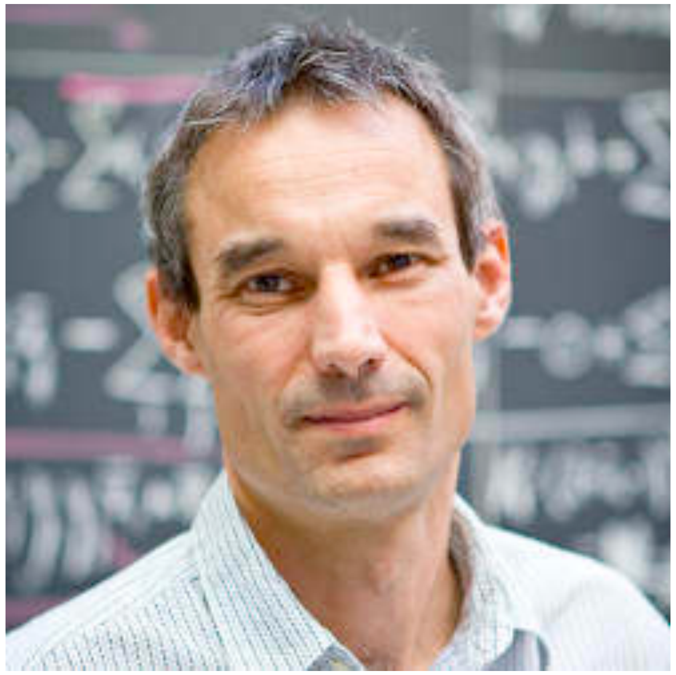

Frans Kaashoek,
Systems

Figure out what you love, and do it (since you are going to be doing it for a while).

If you cannot figure out that there is something that you actually want to do, you will drift, lose your way and fail.

Dave Clark,
Internet Policy

Srini Devadas, Security

The most important criterion is to "explore (and embrace) things you don't understand, as opposed to avoiding things you don't understand."

Take what you are good at and grow it.

Aude Oliva,
Vision

“They didn’t know it was impossible, so they did it.” Too many students think it isn’t possible to do something, so what stops them is themselves. Once they unlock what was stopping them, nothing is impossible.

hink the most important thing in research
a story -- not a theorem or an algorithm
but the story that makes the theorem or
gorithm interesting and exciting.   It’s
nportant to have an "ear" for a good
ory... when do the stories make sense,
hen are they bogus?

Tomas Lozano-Perez
Robotics

The advisor

The best students are possessed by a problem. They're independent. They teach their advisors. They don't do what they're told...they do something more interesting.

(But maybe that's a scary thing to tell new students....but it's true).

Leslie Kaelbling,
Robotics, Machine Learning

Polina Golland,
medical applications of computer vision

Be more stubborn than your advisor

Don't tell your advisor you're doing what they advised against until you’ve solved the problem.

Manolis Kellis,
Computational Biology

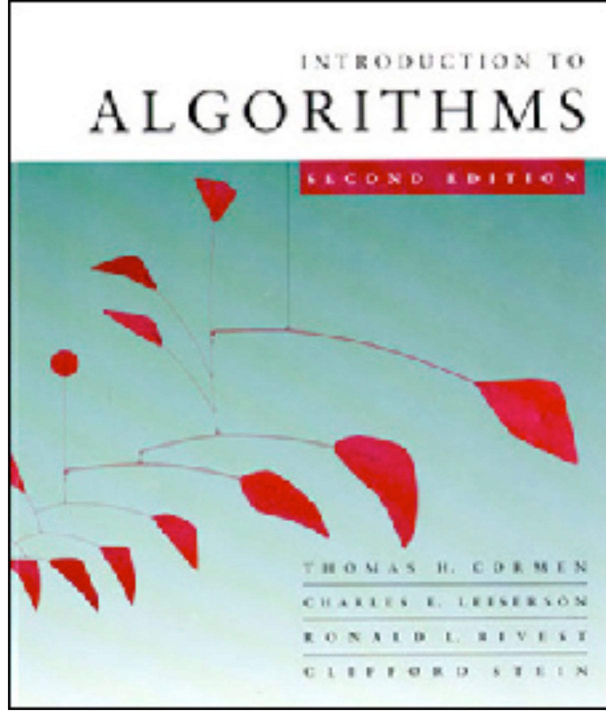

Charles Leiserson,
algorithms

Bill,

I'll tell you a joke instead.

A rabbit is caught by a wolf. The wolf is about to eat the rabbit, but the rabbit protests, "I'm only a few weeks away from defending my Ph.D. dissertation. I have worked so hard, and everyone tells me I have so much promise. It would be a shame to kill me when I have such a bright future of research contributions ahead of me that will benefit the world." The wolf says, "You're writing a Ph.D. dissertation? What is it on?" The rabbit replies, "It's entitled, *The Superiority of Rabbits over Foxes and Wolves*." The wolf says, "That's about the stupidest thing you could have said. I'll eat you right now." The rabbit says, "Wait, wait! Come to my den and read my thesis draft. If you don't agree with my conclusions, I will willingly give myself up to you." So, the wolf goes off with the rabbit to the rabbit's den ... and the wolf never comes out.

A few weeks later, the rabbit is caught by a fox. The fox is about to eat the rabbit, but the rabbit protests, "I'm only a few days away from defending my Ph.D. dissertation. I have worked so hard, and everyone tells me I have so much promise. It would be a shame to kill me when I have such a bright future of research contributions ahead of me that will benefit the world." The fox says, "You're writing a Ph.D. dissertation? What is it on?" The rabbit replies, "It's entitled, *The Superiority of Rabbits over Foxes and Wolves*." The fox says, "That's about the stupidest thing you could have said. I'll eat you right now." The rabbit says, "Wait, wait! Come to my den and read my thesis draft. If you don't agree with my conclusions, I will willingly give myself up to you." So, the fox goes off with the rabbit to the rabbit's den ... and the fox never comes out.

Charles Leiserson

A few weeks later, the rabbit is out and meets up with his old friend the muskrat. The muskrat says, "I hear you finally earned your Ph.D. Congratulations!" The rabbit says, "Yes, I just defended my Ph.D. thesis a few days ago." The muskrat asks, "What was your thesis topic?" The rabbit answers, *The Superiority of Rabbits over Foxes and Wolves*." The muskrat says, "That's quite interesting. Can I read it?" The rabbit says, "Sure. Come to my den." They enter the den, and the muskrat sees the bones of foxes and wolves all over the floor. In the corner is a large lion.

Which brings us to the moral of the story: *More important than your thesis topic is who your advisor is*.

Maybe instead of "quality for success" (which sounds like it's just about the student), might be better to call them "ingredients" of success?

So here's an ingredient: a supportive advisor. Image attached [*next slide*].

Rob Miller,
Human-Computer Interaction

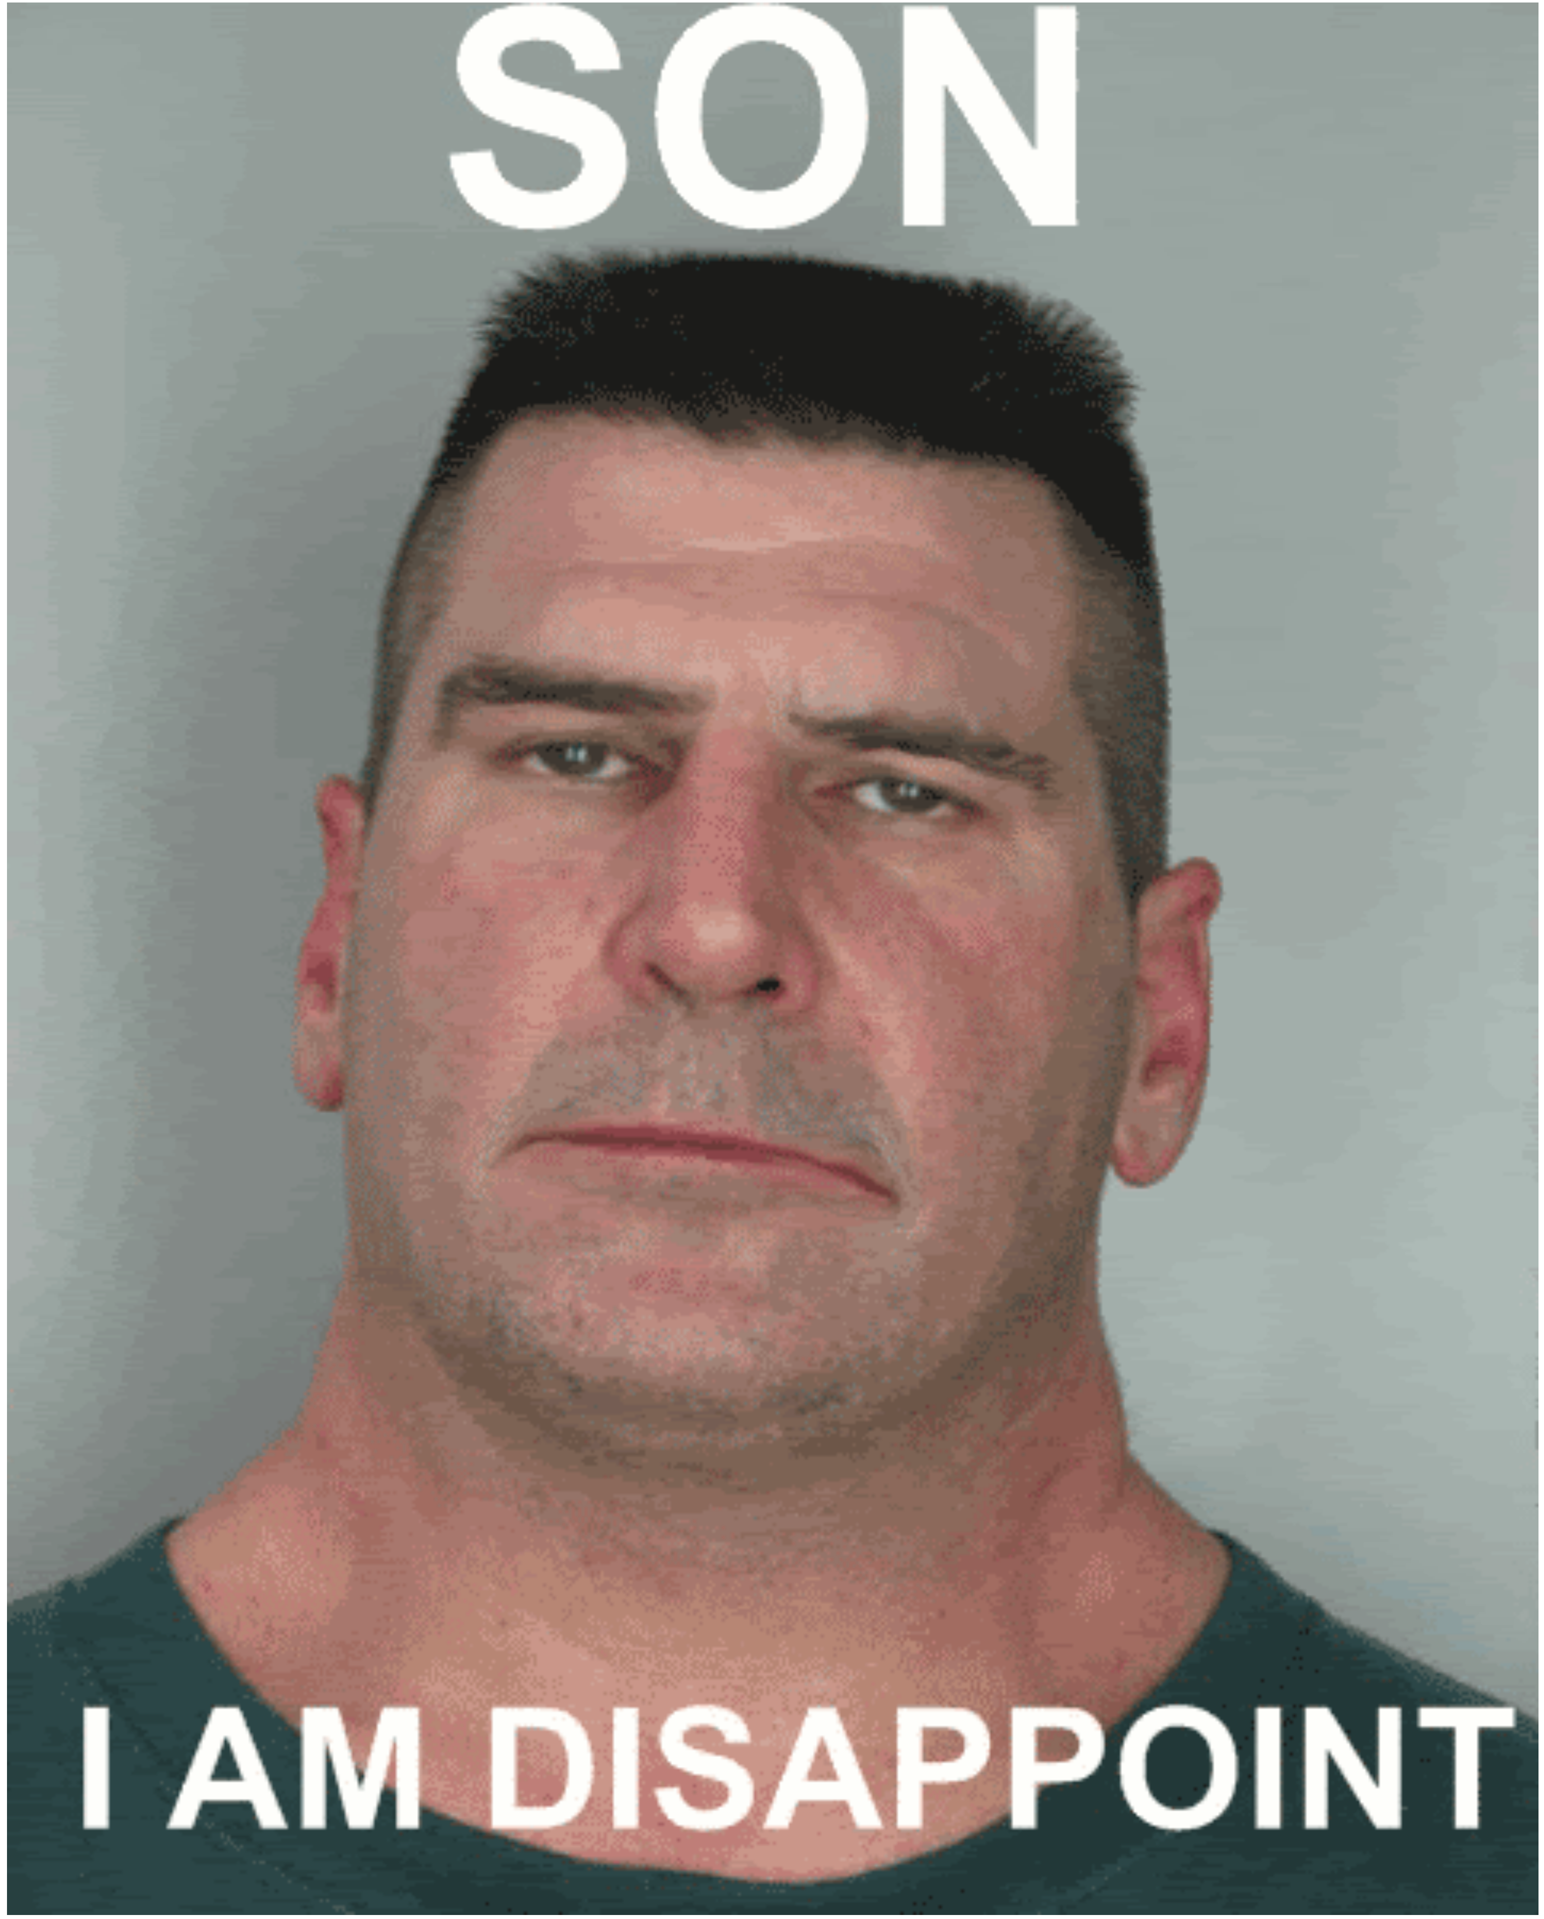
SON
I AM DISAPPOINT

Rules of thumb

Bill, I can offer a couple of rules:

#1: Listen to everything your advisor tells you

#2: Don’t listen to anything your advisor tells you.

there may be a time-dependent alpha in front of rule 1 and (1-alpha) in front of rule 2.
As t -> infinity, alpha -> 0...

#3: (Truly) the most important: believe in yourself and never stop persevering.

Una-May O’Reilly,
Learning

Ruth Rosenholtz,
Vision

* Communicating with your advisor when things are not going well.

* Good grad students need to be Renaissance men/ women: program, experiment, come up with ideas, write, speak well.

* Sometimes you have to cut your losses and work on something else. But otherwise, finish what you start. Persistence, but good judgment.

Daniel Jackson,
Software Design

--Be open: take lots of courses, talk to lots of people, attend seminars.
--Be focused: pick a small project first and get it done.
--Refine your message: at frequent intervals, rework your elevator pitch (and try it out on people who are _not_ in your research area)

Eat, sleep, and breathe a problem until you crack it.

Put everything you have into your problem.

Become the world's foremost expert on your thesis topic.

Surpass your advisor.

Seth Teller,
Robotics, Assistive Technologies

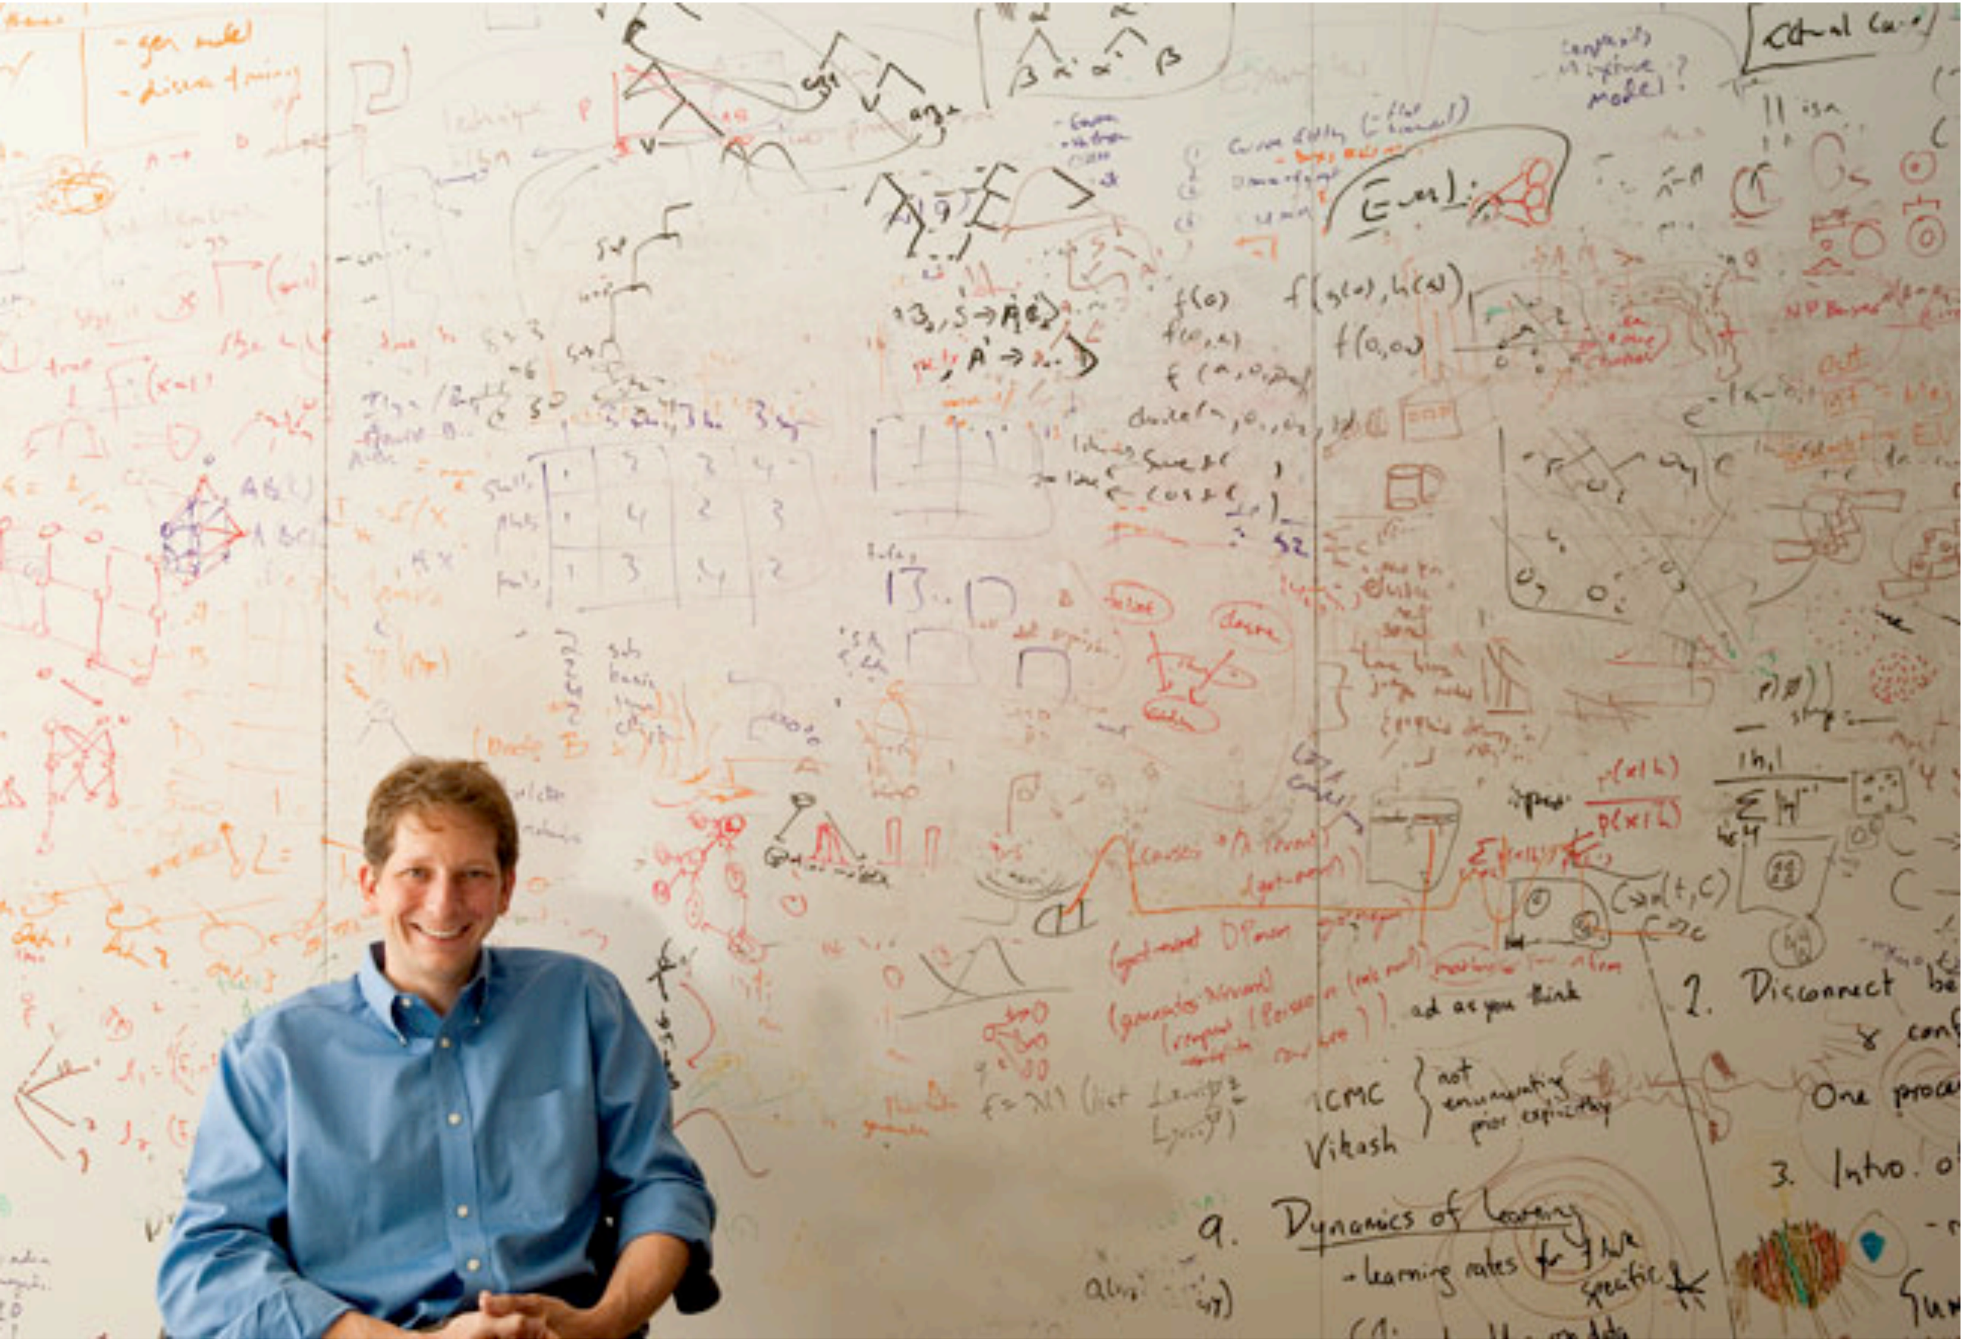

Josh Tenenbaum,
Cognitive Science

(1) Don’t waste time doing research you don’t love.

(2) Don't waste time doing research that other people can do better than you can.

(3) Don't waste time doing research that other people in your field won't care about. (It's okay, and probably a good sign, if some people won't appreciate it, as long as enough people will.)

It's relatively common to find projects that satisfy two out of three of these criteria. But don't settle for your PhD: aim for all three!

Short

Writing and communication.

Taste

Fredo Durand,
Computer Graphics

Being fearless

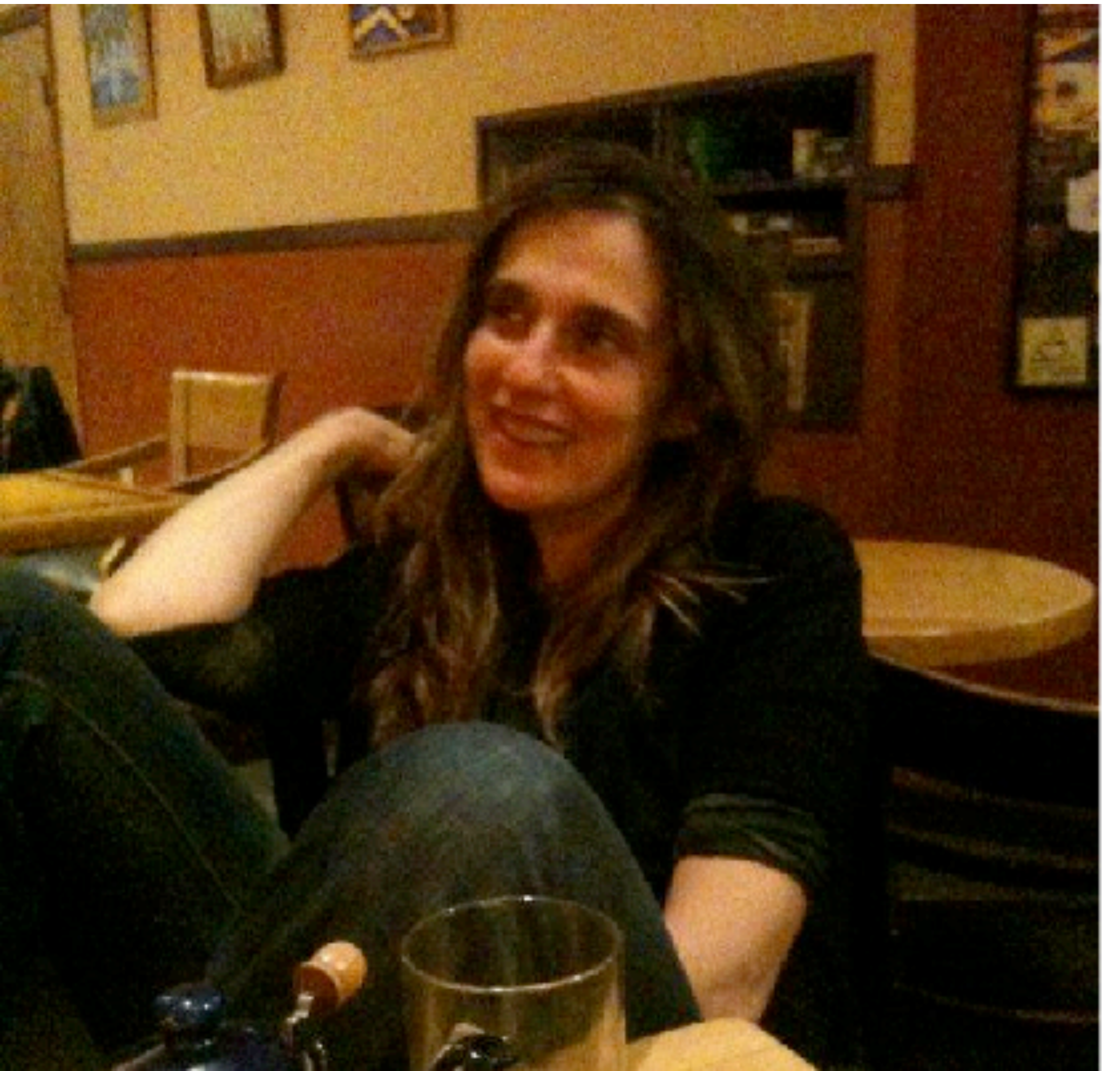

Regina Barzilay,
Natural Language Processing

Passion! A passionate interest in the thesis topic.

Stephanie Seneff,
Spoken Language Processing

Daniela Rus,
Robotics

The passion to pursue an idea despite uncertainties.

Guts!

Silvio Micali,
Theory, Mechanism Design