**Cognitive Vision Group**
University of Maryland Baltimore County
1000 Hilltop Circle, ITE 368, Baltimore MD 21250

# Note to Prospective Ph.D. Students [1]

This short document has been written to help prospective students understand the dynamics of the Cognitive Vision Group. The intention is to summarize the research goals of the group, our philosophy and outlook, our modus operandi, and our values. Hopefully this document gives you a glimpse into the culture and processes of research at CVG. If you are a new member of CVG, this document could also serve as a "handbook" or "starter pack".

## CVG Research Themes

Our mission is to conduct innovative research towards the development of computational theories of visual perception and its links to action, learning, reasoning, and communication.

The current theme of research is "conceptual characterization of visual scenes". Our prior and current work is broadly in the areas of computer vision, machine learning, language grounding, and knowledge representation and reasoning. Current areas of innovation include:

- visual understanding in the absence of complete information;
- adapting and generalizing to new concepts, contexts, domains, environments, and tasks;
- extrapolation across space and time; interpreting and predicting behavior;
- acquisition and communication of visual knowledge from/with/to other agents and humans.

**Coursework / Preparation:** While I don't "require" any specific classes, I expect strong fundamentals (linear algebra, probability, statistics, geometry, calculus, symbolic logic), strong programming skills, and preferably some background in basic engineering such as signal processing, control systems, etc. s Once you start graduate school, I expect rigorous graduate coursework (at or before UMBC) in machine learning and computer vision and some complementary AI knowledge (through courses such as graduate AI, NLP, KRR, RL, etc.). It is unlikely that you will contribute to philosophy (that's what the Ph. in Ph.D. stands for) if you only take classes in computer science – in fact more often than not it leads to a very myopic view of the world. I encourage taking some classes outside the department (eg. genetics, animal behavior, psychology, economics, mathematics, literature).

## CVG Values

We are a group that operates on two primary values: RESPECT and RIGOR.

- **Respect:** We help and support each other to do the best work that we can. This includes participating in discussions, asking questions, giving feedback on papers and presentations, debugging code, celebrating success and navigating failure. We express our opinions respectfully and in good faith. There is no hierarchy in the group – I'm not your boss – I'm an advisor and facilitator in your scientific journey. You shouldn't call me Professor/Dr/Sir/Lord/Supreme Leader etc.
- **Rigor:** The work that we do is rigorous. We are grounded in fundamentals. We pay attention to detail. We are proud of what we do, we are always looking to improve our skills, and we engage with feedback and constructive criticism. We don't gate-keep. We are transparent. We accept our mistakes.

[1] Partly modeled upon similar notes compiled by Shari Liu and Ranjay Krishna

## CVG Meetings

There is a CVG meeting (a "group meeting") every "even week". All CVG members and affiliates are invited which means that the attendees are typically a mix of undergrads, masters, Ph.D, postdocs, and faculty. Typical agenda for group meetings could be sharing results from ongoing experiments and seeking feedback from the attendees; practice talks for accepted papers, proposals, defenses, etc.; discussion of interesting papers; invited talks from visitors; and/or quick updates from everyone.

I meet with each Ph.D. student one-on-one for 30 minutes every "odd week". We can use this time for project discussions (results, updates . . . ), brainstorming (new ideas, getting unstuck, . . . ), editing paper drafts, or sharing new papers or directions that you're excited about. I expect that you'll come prepared. If you have nothing to talk about, but you're available: we can talk about your classes, your goals (see the "Mentor" paragraph above), we could complain about something together, or I might start talking about history and/or language or some other random thing. Sometimes you might be busy implementing and running experiments or we've already spoken recently or you can't make it because of personal reasons – feel free to cancel the meeting.

Often, you could be working on a project with collaborators outside CVG – we will have separate meetings dedicated to such projects. Some students may be co-advised – in such cases we may have a joint meeting with both advisors every now and then.

There are also several seminars, invited talks, and social events that the department/college/university organizes – I encourage you to attend these (free food!).

## Advising Style

One question that I often get from Ph.D. applicants is: *"What is your advising style?"* and the shortest possible answer I can give is *"it depends"*.

Having advised students at different stages and in different capacities, I can confidently say that ain't no such thing as a universal method for advising. My advising style is *different* for each student, depending on what stage the student is in and my mental model of what I think might work for the student. My role as an advisor will evolve the longer we work together and my mental model will get updated the more we communicate about what works and what doesn't.

I have tried to categorize the different roles I expect to play as an advisor:

- **Mentor.** Every semester, I'll ask you to set and share your research goals and "skills" goals. The box below shows my goals during my second semester of graduate school at ASU. (Note: I did not write them down in this format but I wish I had.)

  **Research:** (1) polish and submit a paper I worked on in the first semester `(← this became my first workshop paper  CVPR 2019)` (2) further investigate and build on the preliminary findings from a class project `(← the results weren't great so I abandoned this.)` (3) start investigating whether VQA models understand logic `(← one year later this became my first conference paper  ECCV 2020)`
  **Skills:** (1) become a LaTeX magician (2) work on rapidly converting an idea into an initial action-plan (3) figure out what people mean when they say things like *knowledge* or *reasoning*

  My job as a mentor is to help you make progress towards these goals. I will help you convert goals into actionable plans and serve as a coach. A Ph.D. is tough. It can be frustrating and grueling. There will be a lot of failures. There could be personal challenges as well – if there are ways in which I can understand your situation or just listen, please feel free to share. You're not alone in this journey – I

am dedicated to supporting you, not just your research.

- **Student.** One of the reasons I am in academia as the director of a research group is to learn from my advisees. This is my favorite part of the job. This is also the reason why I prefer working with students whose interests are atleast an $\epsilon$ distance from mine so that we can learn something new together and push the bounds of science. During the course of your Ph.D. you will ask research questions that no one has asked before and you will find evidence to support your hypotheses – when you present such evidence to me (or to anyone reading your papers or watching your presentations), you are the teacher and I am the student. I look forward to learning from you!
- **Teacher.** I typically teach classes in computer vision, machine learning, neural networks, or seminars in advanced topics. If you're at UMBC and are interested in working with me, the best way to get a headstart into my research area would be to take one of my classes. I can teach you things, both in a formal setting through classes and in an informal setting via discussions. For example, I've often found that CS undergrads lack basic engineering education; for instance fundamentals of signal processing, Fourier domain analysis, circuits, and optics – you will either learn this in my computer vision class or during our meetings (there's no escaping this!).
- **Colleague.** I'll treat you as a colleague and you should too. I may have a little more experience in doing research than you but every project is a new exploration and a new challenge that no one has encountered before and it is very likely that I'm as clueless as you are. That is another reason why I'm in academia – it's fun to ask questions that challenge us, to get confused, and to work with others to find answers. Let's collaborate!
- **Driver.** Sometimes (especially in the first year of your Ph.D.) you may be looking for help in formalizing project ideas. I can help you define the problem area, specific project idea, and create a series of subgoals to make progress on. For new students, I recommend working on and executing well-defined projects as early as possible. If you're able to find such well-defined project ideas on your own, great! But if you aren't, that's perfectly normal – ask me to develop a project problem statement for you to work on! This will help you improve your own problem formulation skills and help you become independent for future projects.
- **Facilitator/Promoter.** I will help you write impactful papers, develop engaging talks, find appropriate collaborators, build a network, and find opportunities to grow in your career.

At the end of every semester, I usually do a short chat with each student to understand what's working and what's not to help me improve and finetune my advising for you. But you don't need to wait till the end of the semester – if you have suggestions for me, share them as soon as possible!

### Research Philosophy

Another question I'm often asked is: *"What type of research problems should I work on?"* and honestly, it's hard to answer this question because it's pretty loaded. I prefer to work on projects that could potentially have a lasting impact – questions, ideas, or findings that might still be found useful years later. Tinkering and engineering a practical and "better" solution is cool but that can be done in industry (for a lot more money). 'Applied" papers that take the form of *"we applied technique X to problem Y and it worked"* often have quite a lot of utility, but that's something I'm not interested in. It's much harder to ask brand new questions (new problem formulation) than it is to find a better solution for an old problem (benchmark climbing). I do find some papers that propose incrementally better solutions very insightful, I'm more excited about working on projects that ask interesting questions and perform early explorations towards the answer – these are the types of papers that create subfields (and sometimes entire fields) of inquiry.

When the "GAN" idea was developed by Ian Goodfellow and others, was it the best way to learn

representations at the time? Not really. Did it bring about a paradigm shift in generative modeling? You bet!

Long story short, be ambitious. Propose ideas that you think might have potential to be adopted by many others in the future and set up the foundation for your research philosophy. Remember, this is a doctorate in philosophy.

### Good Practices

A big chunk of the so-called passion or drive comes from within. But lack of organization/time management and environmental factors can also kill passion. Here are some good practices that may help keep you excited to come to the lab and work towards your goal:

- **Project Tracking.** One Doc for One Project. When working on a project, just create a simple Google Doc and share it with all collaborators. The first page of the doc should have a *key idea / summary / abstract* of the project (which you can keep updating as you make progress). Keep adding descriptions of your hypothesis, methods/pipelines, experimental setup, results, analyses, notes, meeting summaries, to-dos in the same doc. This is nice and easy because it doesn't take much of an effort to dump updates to the doc and there's one place for viewing the "history" or log of the project. The single doc helps me and other collaborators quickly recall salient information about the project before meetings and this saves everyone a LOT of time. It also helps you because when it's time to write the paper, you don't have to remember details from months of work – you can just refer to the doc!
- **Read!** Read a lot of papers, but don't overwhelm yourself. Practice reading hygiene – one way to do this is to allocate a fixed amount of time (eg. 1 hr) daily for reading one paper. Don't limit your reading to CS or AI – read so that you have a well-rounded worldview and a unique perspective. Just because something is "hot" or "trending" doesn't mean you need to chase it.
- **Write!** You can't write strong impactful papers if you don't like writing. Ideation shouldn't just happen inside your head – developing good ideas (and dumping bad ones) is a lot easier if you write them down or draw them as workflows/mindmaps. Write stuff down in private documents, however bad you think the ideas may be.
- **Form a Community!** Share ideas, learn from others, promote each other, ask for help and offer help, organize fun activities / social events, seek collaborations (in UMBC and outside).

Best of luck!

**Tejas Gokhale**

Assistant Professor & Director of the Cognitive Vision Group

Computer Science and Electrical Engineering
University of Maryland Baltimore County

Website: https://www.tejasgokhale.com/
Email: gokhale@umbc.edu