

Literature Review Assignment: Interdisciplinary Research Methods

Chirayu Salgarkar

September 19, 2024

Problem #1. I am using BibTeX as my citation environment manager.

Problem #2. I have set up a Public Google Scholar Account. Here is a screenshot.

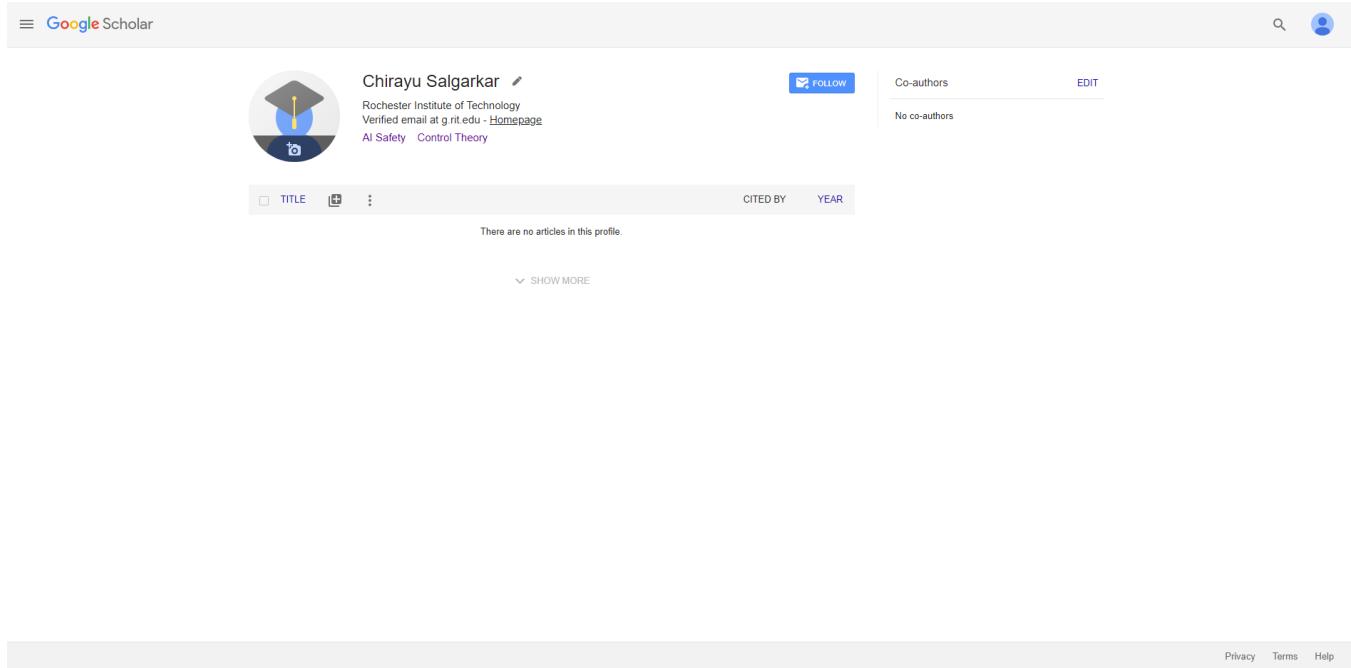


Figure 1: Screenshot of Google Scholar Account. I have set it up as public, but since there are no publications at the outset, searchability is not good.

Problem #2. I have set up Google Scholar Alerts. See Figure 2 for details.

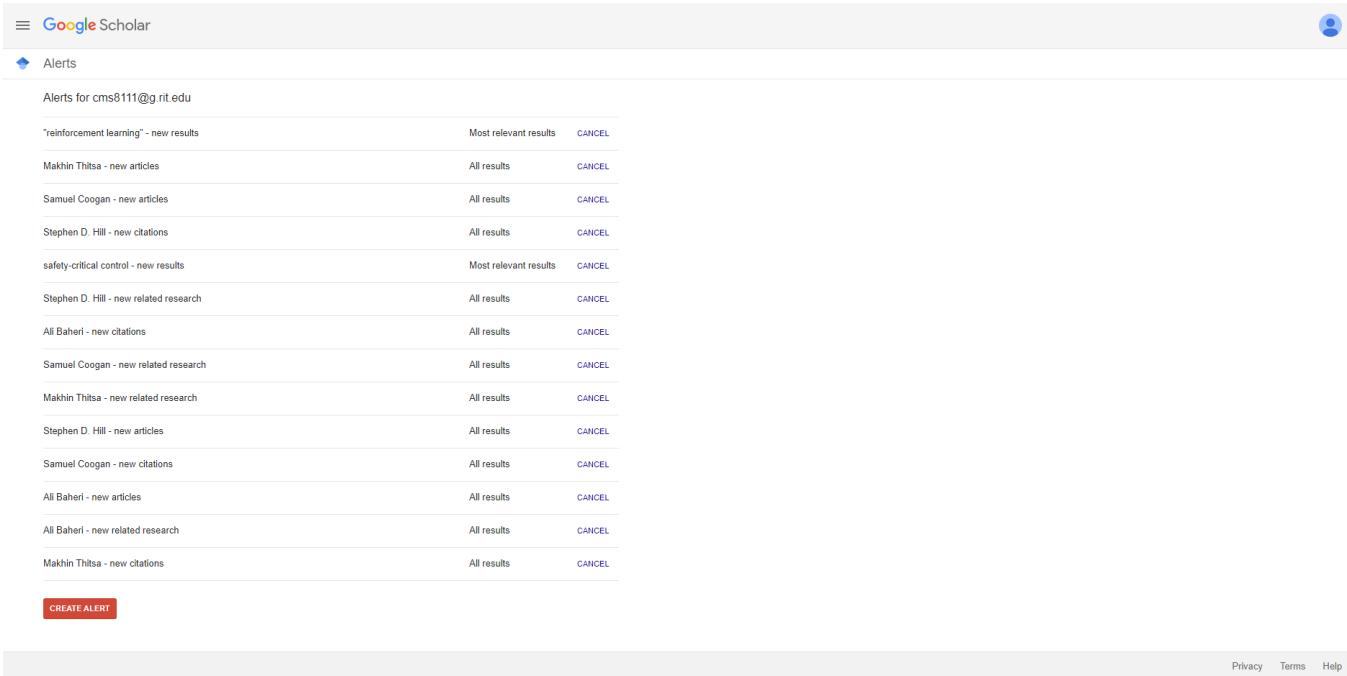


Figure 2: Screenshot of Google Scholar Alerts. Note that it includes both researchers and interesting keywords.

Problem #3. Here are three journals:

1. IEEE Open Journal of Control Systems
2. Automatica
3. Journal of Machine Learning Research

See Table 1 for more details.

Table 1: Journal information summary, with all details given.

Journal	Publisher	Publication Frequency	Average Review Time
IEEE Open Journal of Control Systems	IEEE	Continuous, articles released individually	8-12 weeks
Automatica	Elsevier	Monthly, 28 articles in recent issue	358 days (ScienceDirect), typically 9+ months
Journal of Machine Learning Research (JMLR)	JMLR, Inc. & Microtome Publishing	Continuous, 400 articles in 2023 compendium	3-6 months

IEEE Open Journal of Control Systems is published by IEEE. Publication of accepted papers on IEEE Xplore is guaranteed to be no more than 24 weeks after initial submission. This also means that articles are released individually rather than in batches. Review times are generally on the order of 8 to 12 weeks, but this can vary.

Automatica is published by Elsevier, with issues being monthly. The most current issue had 28 articles published. Average review times seem long, with average submission to acceptance times being 358 days according to ScienceDirect, but these may vary. Reviews have stated that time to publish is usually over 9 months.

Journal of Machine Learning Research (JMLR) is published by JMLR, Inc. and Microtome Publishing. Articles are published continuously upon acceptance, rather than in journal blocks, with average review time being about 3 – 6 months. Volume 24 (a compendium of all publications in the journal in 2023) had 400 entries.

Problem #4. Reinforcement learning and control strategies have converged in recent years, such as reinforcement learning based optimal control methods (Shi et al., 2025) or in the emerging field of online non-stochastic control and decision making (Yan, Zhao, & Zhou, 2023). Engineered systems are designed to be safe, and so recent control theoretical work has focused on approaches to provide provable safety guarantees for critical, but high-dimensional systems, such as in Reinforcement Learning (RL). Recently, methods using Hamilton-Jacobi (HJ) reachability analysis have been proposed to provide provable safety guarantees for RL, by learning the HJ reachability value function while learning control policies (Ganai, Gao, & Herbert, 2024). Variations of Control Barrier Functions (CBFs) can also be used for ensuring the safety of learning-enabled systems, particularly for safe exploration for online RL (Cohen & Belta, 2023). Interesting further developments could include implementing provable safety guarantees in policy-guided diffusion, a method to use diffusion models to generate trajectories and controllable synthetic training data (Jackson et al., 2024). There is also promise in incorporating such work in RL paradigms with different notions of safety, such as welfare-centric fair reinforcement-learning, which casts fair learning as welfare-minimizing (non-maleficence) (Cousins, Asadi, & Littman, 2022).

Problem #5. Before I start designing my presentation, I should *know my audience*. That is, I should see if it is meant to be for a conference, or a group meeting, or for a seminar. I should also be wary of *time constraints*, including how much time I will need for questions and answers. Finally, before I start, I should decide on my goal and the main message I'd like to convey.

References

Cohen, M. H., & Belta, C. (2023). Safe exploration in model-based reinforcement learning using control barrier functions. *Automatica*, 147, 110684.

Cousins, C., Asadi, K., & Littman, M. L. (2022). Fair e3: Efficient welfare-centric fair reinforcement learning. In *5th multidisciplinary conference on reinforcement learning and decision making. rldm*.

Ganai, M., Gao, S., & Herbert, S. (2024). Hamilton-jacobi reachability in reinforcement learning: A survey. *IEEE Open Journal of Control Systems*.

Jackson, M. T., Matthews, M. T., Lu, C., Ellis, B., Whiteson, S., & Foerster, J. (2024). Policy-guided diffusion. *arXiv preprint arXiv:2404.06356*.

Shi, X., Li, Y., Du, C., Chen, C., Zong, G., & Gui, W. (2025). Reinforcement learning-based optimal control for markov jump systems with completely unknown dynamics. *Automatica*, 171, 111886.

Yan, Y.-H., Zhao, P., & Zhou, Z.-H. (2023). Online non-stochastic control with partial feedback. *Journal of Machine Learning Research*, 24(273), 1–50.