Research Statement

Zhiyu Zhang

Summary Adaptivity is an increasingly important focus in modern data science: with minimum human intervention, algorithms often need to self-adjust to the environments they are deployed in, while avoiding catastrophic failure modes. **My passion is to design such adaptive methods for rigorous sequential decision making problems, leveraging ideas from optimization, statistics, signal processing and game theory.** The results are novel theoretical frameworks that blur the boundary between different technical fields, as well as principled solutions to a number of practical challenges.

Advances in artificial intelligence are often motivated by the strengths of biological intelligence, and the pursuit of *adaptivity* is one of the most prominent examples. Think about how humans make everyday decisions: with the proper experience (i.e., inductive bias), we are able to independently thrive in complicated changing environments, while simultaneously refining our experience through this process. In contrast, even the state-of-the-art decision making algorithms fall short in comparison, as extensive human inputs are often required throughout their life cycle, and a slight adversarial perturbation can sometimes make them crumble. This raises an important question for the research community:

How to make algorithms more adaptive to the environments they are deployed in?

My research centers around this high level question, but instantiates it in rigorous data science problems at the intersection of optimization, statistics, signal processing and game theory. Specifically, **I design adaptive sequential decision making algorithms in the following concrete manners.**

- Algorithmically, they have fewer hyperparameters than non-adaptive algorithms, and their proper functioning is based on less stringent assumptions (e.g., the noise is not assumed to be *iid*).
- Quantitatively, they are equipped with optimal performance guarantees that scale with the complexity of the *encountered environment* (which is unknown before deploying the algorithm), rather than the complexity of a pre-defined environment class.

Topic	**Key insight**
Parameter-free optimization [ZCP22c, ZCP22b, ZYCP24]	PDEs reveal key optimization structures in continuous time.
Time series forecasting [ZCP23]	Wavelet + parameter-free optimization improves dynamic regret minimizers.
Uncertainty quantification [ZBY24, ZLY24]	"Adversarial Bayes" unifies stochastic & adversarial conformal prediction.

Table 1: Summary of my representative results and their insights.

Achieving such adaptivity requires uncovering the latent connections between disparate technical fields, as summarized in Table 1. But besides the endless joy from this theoretical endeavor, I am still an engineer at heart, and **an important goal of my research is to address practical data science challenges using the derived insights.** This includes adaptive regularizers that mitigate the *loss of plasticity* in continual learning [MZY24], trustworthy confidence set predictors without statistical assumptions [ZLY24], unsupervised adaptation of confidence set predictors [KZYT24], and beyond.

Zooming out to an even higher level, **my research philosophy / taste is to do "elegant and practical theory"**. Selecting the right research topic is particularly important, as a balance needs to be reached between its intellectual depth and practical relevance. I specifically favor topics with simple settings but rich structures, which means that elegant solutions might be built from **nontrivial twists on the conventional wisdom,** with as few technical nuances as possible (as shown later, my prior works bear this out). From a practical perspective, such topics also quite frequently coincide with the right *level of abstraction*: although it seems hopeless to obtain a mechanistic understanding of many practical data science regimes (especially related to modern vision-language models), I still believe that studying the right theoretical abstraction can provide valuable engineering insights, which **leads (rather than follows) the advances in practice.**

Past Research

Regarding the central theme of adaptivity, my representative results are summarized into the three topics below (also see Table 1).

Parameter-Free Optimization / Online Learning My main technical background is Online Convex Optimization (OCO) [Zin03], an iconic problem setting in *online learning* that bridges convex optimization and game theory. Here, the optimization algorithm faces a sequence of convex loss functions selected by an unknown adversary, and the goal is to ensure low *excess total loss* (i.e., low *regret*) with respect to the optimal fixed iterate (i.e., the optimal *comparator*). Viewing the adversary as a generic noise-generating mechanism, such a setting has been widely adopted to study stochastic optimization and its numerous extensions, including but not limited to training machine learning models [DHS11], quantifying their uncertainty [GC21], and controlling dynamical systems [ABH+19].

Within OCO, strong forms of adaptivity have been studied under the notion of *parameter-freeness* [MS12, MO14], and it is known that certain variants of the celebrated *Follow the Regularized Leader* (FTRL) algorithm [OP16, CO18, MK20] can almost match the performance of the optimally-tuned gradient descent without any tuning. The limitation is that the design of such algorithms is more of an art than a science, and the key problem structures are often concealed by the complicated algebraic analysis (which is a quite common obstacle in optimization theory), making the search of *optimal algorithms* difficult.

A main thread of my past research is developing a novel *continuous time* (CT) framework to streamline the analysis and design better algorithms. The intuition follows naturally from comparing OCO with adjacent fields such as signal processing and control systems, where the usual strategy is to first design an algorithm in CT and then discretize it, with the CT problem being much simpler. Developing an analogous argument for OCO requires addressing a number of technical challenges. More specifically,

1. By bridging the core algorithmic reductions in parameter-free OCO [MO14, CO18] with recent CT characterizations of repeated games [DK20, HLPR23], my first work in this thread [ZCP22c, ICML'22] demonstrates how to convert solutions of a particular PDE (called the *Backward Heat Equation*; BHE) into FTRL *regularizers*. Then, with a proper *discretization argument* developed in [ZYCP24, ALT'24], this provides a tractable "continuous" way to design and analyze a larger class of parameter-free OCO algorithms, whose "discrete" algebraic analysis was prohibitively difficult before.

   In particular, this CT framework generates a nonstandard algorithm based on the *imaginary error function*, achieving stronger notions of optimality than prior works.
2. Next, another work of mine [ZCP22b, NeurIPS'22] applies the CT perspective on a variant of the OCO problem called *OCO with switching cost*. This is a fundamental building block of an adversarial theory of control systems [ABH+19], with a complicated, suboptimally adaptive algorithm already proposed in [ZCP22a, AISTATS'22]. Here, the CT framework reveals that adding switching costs corresponds to giving the BHE a larger *negative diffusivity constant*, which ultimately results in a simpler but optimal algorithm.

Overall, this line of works is particularly close to my heart, as it connects a diverse range of technical pieces and turns them into a coherent pipeline to design better adaptive algorithms. Does it actually make a difference in practice? Focusing on the application in *continual reinforcement learning*, my recent work [MZY24, NeurIPS'24] provides an example which confirms that. The idea is that continual learning algorithms need suitable regularizations to maintain *plasticity* (i.e., the ability to respond to changing inputs), and parameter-free OCO, at its core, provides *implicit regularization schemes* that achieve this well.

Time Series Forecasting My second major research topic is the theoretical foundation of time series forecasting, which can be rephrased as a *dynamic* extension of the above OCO problem. The difference is that instead of requiring the comparator to be a fixed point on the domain, the dynamic setting allows it to change over time (think of it as the ground truth time series), and effectively, the regret is defined with respect to such *comparator sequences*. This brings a substantial shift in the intuition: the problem is closer to signal processing than optimization, as the output of reasonable algorithms typically *does not converge*.

Adaptivity is right at the center of this problem, as one would expect the forecasting performance to depend on the *regularity* of the time series, usually measured by its *variability*. However, existing attempts from the pure optimization perspective are somewhat odd [Zin03, HW15, ZLZ18, JC22]: although the algebra works well, there is in general a lack of intuitive transparency in the exploitation of comparator structures, making further improvements difficult. In particular, an important solution concept called "second-order comparator adaptivity" had not been achieved before.

My contribution is an algorithmic framework [ZCP23, NeurIPS'23] bridging the insights from signal processing and parameter-free OCO. It is known separately that

1. (Signal processing) If the ground truth time series has low variability, then it can be represented as a *sparse* linear combination of the *Haar wavelet basis*.
2. (Parameter-free OCO) There exists a class of OCO algorithms whose (static) regret bounds scale with the sparsity of the fixed comparator.

Combining them leads to a natural strategy which is quite different from all previous attempts: applying sparsity-adaptive OCO algorithms on the *transform domain* of the Haar wavelet basis (i.e., the space of weights). More surprising is its quantitative power: the desirable *second-order comparator adaptive* regret bound can be obtained from the state-of-the-art result on parameter-free OCO (my first research direction) and the *wavelet approximation theory*. **This constitutes the backbone of my PhD dissertation [Zha23], which won the departmental outstanding dissertation award.**

Uncertainty Quantification Since starting my postdoc, I also became interested in the topic of *adversarial* uncertainty quantification. This is motivated by a common practice of modern machine learning: instead of predicting a single label, the model often needs to predict a collection of labels with a specified confidence level, with few or no assumption on the data. The key challenge is to align the model's own perception of uncertainty with its actual performance in the real world.

Conformal prediction (CP) [VGS05] has recently emerged as a premier framework to address this challenge, as it blends the empirical strength of modern ML with the theoretical soundness of traditional statistical methods. The conventional idea (called *split CP*) is to assume iid or *exchangeability* on the environment, such that with a standard *quantile estimation* subroutine, the uncertainty of the ML model itself can be reliably calibrated by an offline dataset. What if such statistical assumptions do not hold? A recent trend is using ideas from OCO to tackle certain surrogate objectives (e.g., controlling the *coverage frequency error* [GC21]), but overall, there are still many open problems regarding the validity of such objectives and the corresponding algorithms.

My latest representative results are CP algorithms built on deeper connections between adversarial online learning and (Bayesian) statistics.

1. First, my prior work [ZBY24, ICML'24] develops an adaptive framework for the fully adversarial setting of CP [GC21], achieving accelerated upper bounds on the coverage frequency error. The key observation is that the FTRL-type OCO algorithms are more compatible with the structure of CP (especially with the coverage frequency error as the performance metric), even though from the perspective of regret they often perform similarly as gradient descent.
2. Recently I started to think about an even more intriguing problem: how to unify the strengths of the conventional split CP and more recent OCO-based approaches? This is an important adaptivity issue previously overlooked by the community, as after all, the users would like to avoid artificial statistical assumptions on the environment, and a good algorithm needs to automatically achieve stronger performance guarantees if the environment is actually easy (iid or exchangeable).

   My recent work [ZLY24, arXiv'24] provides a compelling solution, where the key idea is to replace the conventional frequentist quantile estimator in split CP by its Bayesian counterpart. The insight is that without sacrificing performance under the iid assumption, the Bayesian procedure enforces a regularization that "robustifies" split CP in adversarial environments, and quantitatively, the optimal regret bound can be obtained by interpreting the algorithm as a strong form of FTRL (i.e., without loss linearization). This is also intriguingly related to some recent progresses on *loss-agnostic decision making* [KLST23, LSS24].

Future Plan

Going back to a higher level discussion, my tentative future research plan is sketched as follows.

Research Direction Although there are a number of intriguing open problems I would like to work on in the near future, here I focus on longer term directions (3-5 years). All of them are closely related to my main research theme of adaptivity.

1. **Adversarial uncertainty quantification.** This is a topic I have been thinking about for the past year, and it continues to attract me for multiple reasons. Practically it is associated with high impact, since the field of uncertainty quantification is moving towards general use cases beyond traditional statistical assumptions (such as LLM-powered robotics [RDB+23]), and solving that requires transformative insights rather than just "scaling up". Intellectually the problem is also really deep, as the very basic definition of uncertainty and confidence is nontrivial in the adversarial setting, and furthermore, it remains somewhat unclear what makes the uncertainty evaluation "trustworthy". I believe the answer lies in an improved understanding of the interplay between adversarial online learning and statistics, and *game-theoretic probability* [SV19] is likely to help.
2. **Decision making with systematic considerations.** Most existing works on sequential decision making are *egocentric*: only targeting the default performance metric (e.g., regret), it is so far quite clear what is the fundamental limit one could achieve, as well as how to achieve it. Much less studied is a systematic, "societal" perspective: rather than functioning alone, the algorithm needs to interact with other parties such as data providers, competitors and downstream users. Each party has its own considerations (e.g., efficiency, privacy, fairness), and it is important for the decision making algorithm to take those into account, ideally in an adaptive manner, to ensure the whole system operates properly.

   Motivated by their numerous instantiations at the intersection of ML and EconCS, I am increasingly interested in such decision making problems whose objectives are defined by external parties that the algorithm interacts with. A concrete example from my prior works is [ZLY24], which shows that for adversarial conformal prediction, a Bayesian technique can ensure that the confidence sets received by different downstream users do not contradict each other.
3. **Adaptivity meets embodied AI.** Recall that at the beginning, I used the adaptivity of natural intelligence to motivate adaptive data science algorithms. Although the two concepts are usually studied separately in the literature, their gap is getting significantly narrower due to recent advances in *embodied intelligence*. The key reasoning is the following: the theoretical concept of adaptivity says that performance guarantees should depend on the *complexity* of the problem instance, but crucially, such a complexity notion also depends on the prior knowledge of the algorithm – better priors make the problem easier. Since embodied AI is actually enabled by *really good* priors from large scale pre-training, an exciting future direction is using more adaptive algorithms to improve the efficiency and reliability of embodied AI systems. A sample research question could be, "can quadruped robots reliably improve themselves in the real world, using more adaptive fine-tuning algorithms?"

   Personally I see embodied AI as one of the most promising directions for the next decade, and my postdoc experience in a robotics group has prepared me to tackle challenges in this space (e.g., [MZY24]). Although my own group will likely focus on foundations, I plan to actively collaborate with colleagues in order to actually deploy our algorithms to the real world.

Style and Value Finally, as I am about to lead my own research group in the near future, I would like to share some thoughts on the style and values of my group.

1. I see the well-being of group members as the top priority. A warm and supportive group can boost the energy of individuals, which will eventually help the research quality in the long run.
2. I am committed to creating a collaborative and intellectually stimulating environment, such that group members can freely share thoughts and learn from each other.
3. I value depth over breadth. As a young group, we will focus on selected core topics (see above for examples), build our technical strengths, and aim to go deep.
4. Scientific independence is crucial. Rather than just chasing the trend, I would expect my group to mainly work on problems we are genuinely excited about (while still being practically relevant), and aim for results with our own identifying style.

References by Myself with Collaborators

[KZYT24] Kevin Kasa, Zhiyu Zhang, Heng Yang, and Graham W Taylor. Adapting conformal prediction to distribution shifts without labels. *arXiv preprint arXiv:2406.01416*, 2024.

[MZY24] Aneesh Muppidi, Zhiyu Zhang, and Heng Yang. Fast TRAC: A parameter-free optimizer for lifelong reinforcement learning. *Advances in Neural Information Processing Systems* (NeurIPS), 2024.

[ZBY24] Zhiyu Zhang, David Bombara, and Heng Yang. Discounted adaptive online learning: Towards better regularization. In *International Conference on Machine Learning* (ICML), 2024.

[ZCP22a] Zhiyu Zhang, Ashok Cutkosky, and Ioannis Ch Paschalidis. Adversarial tracking control via strongly adaptive online learning with memory. In *International Conference on Artificial Intelligence and Statistics* (AISTATS), 2022.

[ZCP22b] Zhiyu Zhang, Ashok Cutkosky, and Ioannis Ch Paschalidis. Optimal comparator adaptive online learning with switching cost. *Advances in Neural Information Processing Systems* (NeurIPS), 2022.

[ZCP22c] Zhiyu Zhang, Ashok Cutkosky, and Ioannis Ch Paschalidis. PDE-based optimal strategy for unconstrained online learning. In *International Conference on Machine Learning* (ICML), 2022.

[ZCP23] Zhiyu Zhang, Ashok Cutkosky, and Ioannis Ch Paschalidis. Unconstrained dynamic regret via sparse coding. *Advances in Neural Information Processing Systems* (NeurIPS), 2023.

[Zha23] Zhiyu Zhang. *Temporal aspects of adaptive online learning: Continuity and representation*. PhD thesis, Boston University, 2023.

[ZLY24] Zhiyu Zhang, Zhou Lu, and Heng Yang. The benefit of being Bayesian in online conformal prediction. *arXiv preprint arXiv:2410.02561*, 2024.

[ZYCP24] Zhiyu Zhang, Heng Yang, Ashok Cutkosky, and Ioannis Ch Paschalidis. Improving adaptive online learning using refined discretization. In *International Conference on Algorithmic Learning Theory* (ALT), 2024.

Other References

[ABH+19] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.

[CO18] Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Conference on Learning Theory*, pages 1493–1529. PMLR, 2018.

[DHS11] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.

[DK20] Nadejda Drenska and Robert V Kohn. Prediction with expert advice: A PDE perspective. *Journal of Nonlinear Science*, 30(1):137–173, 2020.

[GC21] Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.

[HLPR23] Nicholas JA Harvey, Christopher Liaw, Edwin Perkins, and Sikander Randhawa. Optimal anytime regret with two experts. *Mathematical Statistics and Learning*, 6(1):87–142, 2023.

[HW15] Eric C Hall and Rebecca M Willett. Online convex optimization in dynamic environments. *IEEE Journal of Selected Topics in Signal Processing*, 9(4):647–662, 2015.

[JC22] Andrew Jacobsen and Ashok Cutkosky. Parameter-free mirror descent. In *Conference on Learning Theory*, pages 4160–4211. PMLR, 2022.

[KLST23] Bobby Kleinberg, Renato Paes Leme, Jon Schneider, and Yifeng Teng. U-calibration: Forecasting for an unknown agent. In *Conference on Learning Theory*, pages 5143–5145. PMLR, 2023.

[LSS24] Haipeng Luo, Spandan Senapati, and Vatsal Sharan. Optimal multiclass U-calibration error and beyond. *Advances in Neural Information Processing Systems*, 2024.

[MK20] Zakaria Mhammedi and Wouter M Koolen. Lipschitz and comparator-norm adaptivity in online learning. In *Conference on Learning Theory*, pages 2858–2887. PMLR, 2020.

[MO14] H Brendan McMahan and Francesco Orabona. Unconstrained online linear learning in Hilbert spaces: Minimax algorithms and normal approximations. In *Conference on Learning Theory*, pages 1020–1039. PMLR, 2014.

[MS12] Brendan Mcmahan and Matthew Streeter. No-regret algorithms for unconstrained online convex optimization. *Advances in Neural Information Processing Systems*, 25, 2012.

[OP16] Francesco Orabona and Dávid Pál. Coin betting and parameter-free online learning. *Advances in Neural Information Processing Systems*, 29, 2016.

[RDB+23] Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, et al. Robots that ask for help: Uncertainty alignment for large language model planners. In *Conference on Robot Learning*, pages 661–682. PMLR, 2023.

[SV19] Glenn Shafer and Vladimir Vovk. *Game-theoretic foundations for probability and finance*, volume 455. John Wiley & Sons, 2019.

[VGS05] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.

[Zin03] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning*, pages 928–936, 2003.

[ZLZ18] Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. *Advances in Neural Information Processing Systems*, 31, 2018.