Research Statement

Arif Dönmez

My research interests span symmetry-informed machine learning, information geometry (Amari, 2016), differentiable programming (Blondel & Roulet, 2024), and latent cause theory (Gershman & Niv, 2012).

It began with my doctoral thesis, in which I focused on classification problems, particularly regarding a specific class of linear algebra problems (16G10[1]) that may not be fully solvable by writing down all possible classifications. I approached the classification problem with methods from algebraic geometry (14D20[2]) (Mumford et al., 1994). I applied geometric-invariant theory for this special type of matrix-problems (Dönmez & Reineke, 2024). After submitting my doctoral thesis, I wanted to migrate the rigorousness of geometric-invariant theory to machine learning, which was mainly motivated to resolve the 'black box' areas of Deep Learning approaches. In mathematics, the study of group actions arises as the object to study to tackle classification problems. And explainability is given by design when studying group actions. Surprisingly, neither groups nor geometric ideas had found their way into machine learning at that time. I discussed my dissatisfaction with the lack of these concepts in machine learning with my postdoc advisors, Markus Reineke, and Axel Mosig, back in 2020, and they motivated me further in my endeavors.

Symmetry-informed machine learning

Mainly inspired by Emmy Noether's work (Noether, 1983b), the stance taken here can be expressed as follows: 'symmetry is the only constant amidst change'. In other words, studying 'what stays invariant' for specific classification problems is fundamental to learning. This is actually what invariant theory aims to study: invariant functions regarding group actions. So, I approached machine learning via symmetries. Unfortunately, I realized naively migrating symmetries to machine learning problems could not succeed. For example, consider the scaled and translated image recognition problem, aiming to cluster images originating from the same image after scaling and translating. Then 'border effects' can occur; an object within the original image is no longer visible in the resulting image after scaling and translating. In such cases, we cannot reconstruct the original image from that extract; no (natural) transformation can reproduce the original image out of the extract. In other words, no group transformation can model this problem. So, the classical group action sense symmetries are not broad enough. At this time, faced with this challenge, Bronstein's work (Bronstein et al., 2021) on the significance of group actions in the context of deep learning was published and attracted considerable attention. I aimed to address the mentioned limitations by taking a rigorous algebraic approach powered by my background. My approach relies on principles from Alexander Grothendieck's 'Rising Sea' philosophy (McLarty, 2007). The long and short of it: 'If a phenomenon seems hard to explain, it's because you haven't fully understood how general it is. Once you figure out how general it is, the explanation will stare you in the face.' One thing that always excited me about Grothendieck's geometry revolution was that he rethought (and eventually redefined) the very notion of a point in geometry. Strongly inspired and influenced by this, I have rethought the classification problem by rigorously generalizing symmetries to partial symmetries of groupoids in my theoretical paper (Dönmez, 2023) (accepted at NeurIPS 2022, NeurReps track).

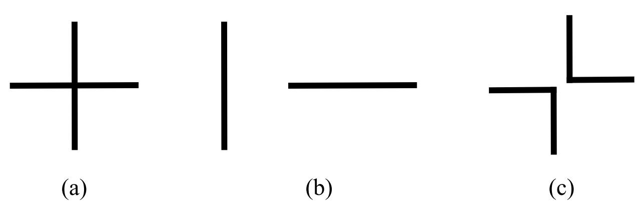

Figure 1. Law of continuity: Lines are always seen as following the simplest path. If two lines cross (a), we do not assume that the course of the lines makes a bend at this point (c), but we see two straight, continuous lines (b).

Another problem in naively migrating the symmetry concept lies in the finite representation capabilities of machines. However, you can find finite symmetries in everything (Noether, 1983a). We have to restrict (partial) symmetries meaningfully. Fundamental, therefore, was the insight from Gestalt psychology that *the whole is more than the sum of its parts* (Koffka, 1935; Köhler, 1967). So, in my follow-up paper (Dönmez, 2024) (accepted at NeurIPS 2023, NeurReps track), I proposed that latent causes are manifestations

[1] 2020 Mathematics Subject Classification: Representations of associative Artinian rings

[2] 2020 Mathematics Subject Classification: Algebraic moduli problems, moduli of vector bundles

of context-dependent symmetries and to restrict symmetries meaningfully, the need for intuitive physics (Wellman & Gelman, 1992; Lake et al., 2017) and Gestalt psychology (law of continuity) on an abstract level (Figure 1). My approach uses linear representations of the symmetric group, statistics, and symbolic computations. I demonstrated in the first application how this approach can replicate experimental outcomes from the cognitive science domain. I plan to integrate further principles from Gestalt psychology and cognitive sciences into machine learning to boost the reasoning capabilities of models.

THE SYMMETRY-LOSS

My follow-up paper, 'On Symmetry-Informed Machine Learning', is under review. It addresses the limitations of my previous work by introducing *symmetry-loss* to enhance applicability. Groups (and groupoids) are fundamental to machine learning, but their combinatorial nature makes them impractical for real-world applications. So, I suggested that instead of identifying a specific group in the data, the choice of groups and their actions should be an inductive bias set by the modeler. This perspective aligns with the principle of free energy (Friston et al., 2006; Buckley et al., 2017), as well as the concept of 'startup software' (Lake et al., 2017). I proposed a differentiable parameterized program with a training objective based on symmetry-loss, supported by invariant theory. One of the key advantages of this approach is that it is inherently explainable.

I explored symmetries in molecular graphs using a symmetry-informed framework to predict the *in vitro* activity of chemicals against transthyretin (TTR), utilizing data from the Tox24 challenge (Tetko, 2024). My model ranked in the top 5% of submissions, even without domain-specific prior knowledge. It effectively explains predictions by weighting chemical bonds based on their activity against TTR, recognizing the importance of hydrocarbon sulfonates and hydrofluorocarbons, which align with recent wet lab findings (Gong et al., 2024; Langberg et al., 2024). My model's ability to replicate these insights suggests strong generalization capabilities, though further verification is needed. I plan to further investigate this topic, especially its effectiveness in various applications, extend it with Bayesian inference, and use it as a generative model. I'm also interested in exploring this promising framework for imaging processes and sequencing data, which can lead to models with significantly fewer parameters. The implication is clear: a better choice of group action results in needing fewer parameters.

Differentiable programming

I recognized the benefits of differentiable programming when introducing the symmetry-informed framework. Similar to how convolutional neural networks have advanced computer vision, differentiable programming can enhance solutions across various domains. It allows for the creation of more explainable models while reducing complexity.

As a follow-up of my performed evaluation and analyses in Galanjuk et al. (2022), I aim to quantify the force generated by beating cardiomyocytes by tracking flow dynamics in the videos. Therefore, I developed an explainable model for detecting laminar flow, motivated by the convolution of probability densities. Our final model, which utilizes differentiable programming, includes fewer than 15,000 parameters, making it suitable for cost-effective Internet of Things (IoT) devices. I plan to combine it with the Hagen–Poiseuille equation for applications in toxicity testing and surveillance systems.

In addition, I want to introduce a differentiable colocalization analysis for evaluating biological microscopy images.

Information geometry for life sciences

Since starting at the Leibniz Institute for Environmental Medicine, I have been working in biostatistics under the supervision of Martin Scholze. During my time in the group, I developed, implemented, and established a pipeline for biostatistical (non-linear) regression analyses of concentration-response data, particularly focusing on the estimation of benchmark levels and benchmark concentrations, which are commonly referred to as effective doses in *in-vivo* experiments. The biostatistical analyses presented in the papers Cediel-Ulloa et al. (2025); Hartmann et al. (2023); Bartmann et al. (2023); Blum et al. (2023); Galanjuk et al. (2022) were produced by me.

After working on these classical biostatistical applications, I shifted my focus to evaluating and expressing uncertainty in hazard characterization. By studying the WHO 'Guidance Document on Evaluating and Expressing Uncertainty in Hazard Characterization' (Organization et al., 2018), I realized that the distributions involved arise from the normal distribution across different Riemannian manifolds. This insight highlighted the necessity of integrating information geometry for more advanced and accurate data inference. My position on this topic is that for applicable statistics on various geometric domains, the heat equation on Riemannian manifolds must be considered.

Working with Martin Parparella, I applied these geometrically-informed statistics to evaluate uncertainty in *in-vitro* hazard characterization at the population level, utilizing statistics within two-dimensional hyperbolic geometry. I aim to expand my R implementation into a comprehensive library that allows users to easily perform statistical analyses on different Riemannian manifolds.

References

Amari, S.-i. *Information geometry and its applications*, volume 194. Springer, 2016.

Bartmann, K., Bendt, F., Dönmez, A., Haag, D., Keßel, H. E., Masjosthusmann, S., Noel, C., Wu, J., Zhou, P., and Fritsche, E. A human ipsc-based in vitro neural network formation assay to investigate neurodevelopmental toxicity of pesticides. *ALTEX-Alternatives to animal experimentation*, 40(3):452–470, 2023.

Blondel, M. and Roulet, V. The elements of differentiable programming. *arXiv preprint arXiv:2403.14606*, 2024.

Blum, J., Masjosthusmann, S., Bartmann, K., Bendt, F., Dolde, X., Dönmez, A., Förster, N., Holzer, A.-K., Hübenthal, U., Keßel, H. E., et al. Establishment of a human cell-based in vitro battery to assess developmental neurotoxicity hazard of chemicals. *Chemosphere*, 311:137035, 2023.

Bronstein, M. M., Bruna, J., Cohen, T., and Veličković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Buckley, C. L., Kim, C. S., McGregor, S., and Seth, A. K. The free energy principle for action and perception: A mathematical review. *Journal of mathematical psychology*, 81:55–79, 2017.

Cediel-Ulloa, A., Awoga, R., Dönmez, A., Yu, X., Gliga, A., Attoff, K., Forsby, A., and Rüegg, J. Characterization of the c17. 2 cell line as test system for endocrine disruption-induced developmental neurotoxicity. *ALTEX-Alternatives to animal experimentation*, 42(1):91–110, 2025.

Dönmez, A. On the ambiguity in classification. In Sanborn, S., Shewmake, C., Azeglio, S., Di Bernardo, A., and Miolane, N. (eds.), *Proceedings of the 1st NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, volume 197 of *Proceedings of Machine Learning Research*, pp. 158–170. PMLR, 03 Dec 2023. URL `https://proceedings.mlr.press/v197/donmez23a.html`.

Dönmez, A. Discovering latent causes and memory modification: A computational approach using symmetry and geometry. In Sanborn, S., Shewmake, C., Azeglio, S., and Miolane, N. (eds.), *Proceedings of the 2nd NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, volume 228 of *Proceedings of Machine Learning Research*, pp. 443–458. PMLR, 16 Dec 2024. URL `https://proceedings.mlr.press/v228/donmez24a.html`.

Dönmez, A. and Reineke, M. Moduli of representations of one-point extensions. *Communications in Algebra*, pp. 1–24, 2024.

Friston, K., Kilner, J., and Harrison, L. A free energy principle for the brain. *Journal of physiology-Paris*, 100(1-3):70–87, 2006.

Galanjuk, S., Zühr, E., Dönmez, A., Bartsch, D., Kurian, L., Tigges, J., and Fritsche, E. The human induced pluripotent stem cell test as an alternative method for embryotoxicity testing. *International Journal of Molecular Sciences*, 23(6):3295, 2022.

Gershman, S. J. and Niv, Y. Exploring a latent cause theory of classical conditioning. *Learning & behavior*, 40:255–268, 2012.

Gong, Y., Sun, J., Wang, X., Barrett, H., and Peng, H. Identification of hydrocarbon sulfonates as previously overlooked transthyretin ligands in the environment. *Environmental Science & Technology*, 58(23):10227–10239, 2024.

Hartmann, J., Henschel, N., Bartmann, K., Dönmez, A., Brockerhoff, G., Koch, K., and Fritsche, E. Molecular and functional characterization of different brainsphere models for use in neurotoxicity testing on microelectrode arrays. *Cells*, 12(9):1270, 2023.

Koffka, K. Principles of gestalt psychology. 1935.

Köhler, W. Gestalt psychology. *Psychologische Forschung*, 31(1): XVIII–XXX, 1967.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.

Langberg, H. A., Choyke, S., Hale, S. E., Koekkoek, J., Cenijn, P. H., Lamoree, M. H., Rundberget, T., Jartun, M., Breedveld, G. D., Jenssen, B. M., et al. Effect-directed analysis based on transthyretin binding activity of per-and polyfluoroalkyl substances in a contaminated sediment extract. *Environmental toxicology and chemistry*, 43(2):245–258, 2024.

McLarty, C. *The rising sea: Grothendieck on simplicity and generality*. na, 2007.

Mumford, D., Fogarty, J., and Kirwan, F. *Geometric invariant theory*, volume 34. Springer Science & Business Media, 1994.

Noether, E. Der endlichkeitssatz der invarianten endlicher gruppen. In *Gesammelte Abhandlungen-Collected Papers*, pp. 141–144. Springer, 1983a.

Noether, E. Invariante variationsprobleme. In *Gesammelte Abhandlungen-Collected Papers*, pp. 231–239. Springer, 1983b.

Organization, W. H. et al. Guidance document on evaluating and expressing uncertainty in hazard characterization. In *Guidance document on evaluating and expressing uncertainty in hazard characterization*. 2018.

Tetko, I. V. Tox24 challenge. *Chemical Research in Toxicology*, 2024.

Wellman, H. M. and Gelman, S. A. Cognitive development: Foundational theories of core domains. *Annual review of psychology*, 43(1):337–375, 1992.