

Skills and expertise

Research interests	Topological phases and topological quantum computing; deep and self-supervised learning for physics
Software development	Python, modern C++, Julia, Linux, \LaTeX , git, jupyter
Frameworks	PyTorch, Qiskit, Hugging Face Transformers, PyTorch3D, TensorFlow
Deep learning	Since 2017, co-authored 20 papers and preprints in multiple fields of deep learning, including top-tier venues (CVPR, JMLR). Wide knowledge of current trends in deep learning.

Education

- 2022 – present **PhD in Physics**, *Technion – Israel Institute of Technology*, Haifa.
- Thesis topic: “Topological Quantum Computing Beyond Majorana Fermions”, advised by Prof. [Netanel Lindner](#);
 - Research in theoretical condensed matter: strongly correlated phases, interfaces in 2D systems (fractional quantum Hall, Kitaev spin liquid, $p + ip$ superconductors);
 - Adams fellow;
 - Teaching experience: “Solid State Physics”.
- 2020 – 2021 **MSc in Computer Science**, *Technion – Israel Institute of Technology*, Haifa, Cum Laude.
- Thesis: “Reducing Supervision in Visual Recognition Tasks”, advised by Prof. [Alex Bronstein](#), Prof. Avi Mendelson, and Dr. Chaim Baskin;
 - Teaching experience: “Advanced Topics in Deep Learning”, “Deep Learning on Computational Accelerators”, “Intro to Machine Learning”, Deep Learning seminar organization;
 - Advising experience: advised research projects on computer vision;
 - Reviewer for T-PAMI, CVPR, ICCV, ECCV, WACV;
 - CS dean excellence scholarship recipient.
- 2016 – 2020 **BSc in Computer Science and BSc in Physics and Mathematics**, *Technion – Israel Institute of Technology*, Haifa, GPA *92.00*, Cum Laude.
- Participant of Rothschild Technion Program for Excellence;
 - ICPC semifinals (SWERC): 2018 – honorable mention, 2019 – bronze medal (11th place).
- Summer 2023 **Princeton Summer School on Condensed Matter Physics**, Princeton.
- Summer 2022 **Topological Matter School**, San Sebastian.
- Summer 2018 **DeepBayes**, *Summer school on Bayesian methods in deep learning*.

Projects and open source contribution

- 2022 **QHack 2022 Hackathon**, “*Barren plateau inhabitants*”, 2nd place at IBM Qiskit Challenge, 1st place Google Quantum AI Research Challenge.
Simulation of anyons within the toric code model.
- 2019 – 2020 **TensorFlow**.
Implemented differentiable eigendecomposition of general matrices for TensorFlow.
- 2016 – 2018 **tiny-dnn**.
Maintainer of tiny-dnn: header only, dependency-free deep learning framework in C++14.

Industrial Experience

- Fall 2020 **Research Intern**, *Snap Research*, Los Angeles (remote), Creative vision group.
- Hosts: Sergey Tulyakov and Olly Woodford;
 - Researched 3D shape reconstruction by training on the dataset of single 2D views;
 - Implemented systems for dense and sparse 3D shape reconstruction from scratch with PyTorch3D.
- 2016 – 2020 **Research Assistant**, *Technion*, Haifa, Professor [Alex Bronstein](#)’s group.
- Investigated compression methods and their impact on DNN performance;
 - Implemented and reproduced the latest DL algorithms and papers.
- Summer 2017 **Google Summer of Code Participant**, *OpenCV*.
GPU enabled deep learning framework: introducing GPU support for [tiny-dnn](#), C++14 header-only deep learning library

Publications

- [1] Moshe Kimhi, Shai Kimhi, **Evgenii Zheltonozhskii**, Or Litany, and Chaim Baskin. "Semi-Supervised Semantic Segmentation via Marginal Contextual Information". In: *arXiv pre-print* (Aug. 2023). arXiv: 2308.13900 [cs.CV]. URL: <https://arxiv.org/abs/2308.13900>.
- [2] Raymond Li et al. "StarCoder: may the source be with you!" In: *arXiv pre-print* (May 2023). arXiv: 2305.06161 [cs.CL].
- [3] Tom Avrech, **Evgenii Zheltonozhskii**, Chaim Baskin, and Ehud Rivlin. "GoToNet: Fast Monocular Scene Exposure and Exploration". In: *Journal of Intelligent & Robotic Systems* 105.3 (July 2022), p. 65. DOI: 10.1007/s10846-022-01646-9. URL: <https://doi.org/10.1007/s10846-022-01646-9>.
- [4] Aarohi Srivastava et al. "Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models". In: *Transactions on Machine Learning Research* (Apr. 2023). ISSN: 2835-8856. URL: <https://openreview.net/forum?id=uyTL5Bvosj>.
- [5] Maxim Fishman, Chaim Baskin, **Evgenii Zheltonozhskii**, Ron Banner, and Avi Mendelson. "On Recoverability of Graph Neural Network Representations". In: *arXiv pre-print* (Jan. 2022). URL: <https://arxiv.org/abs/2201.12843>.
- [6] Adam Botach, **Evgenii Zheltonozhskii**, and Chaim Baskin. "End-to-End Referring Video Object Segmentation with Multimodal Transformers". In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2022. URL: https://openaccess.thecvf.com/content/CVPR2022/html/Botach_End-to-End_Referring_Video_Object_Segmentation_With_Multimodal_Transformers_CVPR_2022_paper.html.
- [7] **Evgenii Zheltonozhskii**, Chaim Baskin, Avi Mendelson, Alex M. Bronstein, and Or Litany. "Contrast to Divide: Self-Supervised Pre-Training for Learning with Noisy Labels". In: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. Jan. 2022, pp. 1657–1667. URL: https://openaccess.thecvf.com/content/WACV2022/html/Zheltonozhskii_Contrast_To_Divide_Self-Supervised_Pre-Training_for_Learning_With_Noisy_Labels_WACV_2022_paper.html.
- [8] Ameen Ali, Tomer Galanti, **Evgenii Zheltonozhskii**, Chaim Baskin, and Lior Wolf. "Weakly Supervised Recovery of Semantic Attributes". In: *First Conference on Causal Learning and Reasoning*. Apr. 2022. URL: <https://openreview.net/forum?id=GdAzRedTV7J>.
- [9] Ben Finkelshtein, Chaim Baskin, **Evgenii Zheltonozhskii**, and Uri Alon. "Single-node attacks for fooling graph neural networks". In: *Neurocomputing* 513 (Nov. 2022), pp. 1–12. ISSN: 0925-2312. DOI: <https://doi.org/10.1016/j.neucom.2022.09.115>. URL: <https://www.sciencedirect.com/science/article/pii/S0925231222012012>.
- [10] **Evgenii Zheltonozhskii**, Chaim Baskin, Alex M. Bronstein, and Avi Mendelson. "Self-Supervised Learning for Large-Scale Unsupervised Image Clustering". In: *NeurIPS Self-Supervised Learning Workshop* (Aug. 2020). URL: <https://arxiv.org/abs/2008.10312>.
- [11] Alex Karbachevsky, Chaim Baskin, **Evgenii Zheltonozhskii***, Yevgeny Yermolin, Freddy Gabbay, Alex M. Bronstein, and Avi Mendelson. "Early-Stage Neural Network Hardware Performance Analysis". In: *Sustainability* 13.2 (Jan. 2021): *Energy-Efficient Computing Systems for Deep Learning*. Ed. by José Cano, José L. Abellán, and David Kaeli, p. 717. ISSN: 2071-1050. DOI: 10.3390/su13020717. URL: <http://dx.doi.org/10.3390/su13020717>.
- [12] **Evgenii Zheltonozhskii**, Chaim Baskin, Yaniv Nemcovsky, Brian Chmiel, Avi Mendelson, and Alex M. Bronstein. "Colored Noise Injection for Training Adversarially Robust Neural Networks". In: *arXiv pre-print* (Mar. 2020). URL: <https://arxiv.org/abs/2003.02188>.
- [13] Yaniv Nemcovsky, **Evgenii Zheltonozhskii***, Chaim Baskin, Brian Chmiel, Alex M. Bronstein, and Avi Mendelson. "Adversarial robustness via noise injection in smoothed models". In: *Applied Intelligence* (Aug. 2022). DOI: 10.1007/s10489-022-03423-5. URL: <https://doi.org/10.1007/s10489-022-03423-5>.
- [14] Yury Nahshan, Brian Chmiel, Chaim Baskin, **Evgenii Zheltonozhskii**, Ron Banner, Alex M. Bronstein, and Avi Mendelson. "Loss Aware Post-Training Quantization". In: *Machine Learning* (Oct. 2021). ISSN: 1573-0565. DOI: 10.1007/s10994-021-06053-z. URL: <https://link.springer.com/article/10.1007/s10994-021-06053-z>.
- [15] Chaim Baskin, Brian Chmiel, **Evgenii Zheltonozhskii***, Ron Banner, Alex M. Bronstein, and Avi Mendelson. "CAT: Compression-Aware Training for Bandwidth Reduction". In: *Journal of Machine Learning Research* 22.269 (Aug. 2021), pp. 1–20. URL: <http://jmlr.org/papers/v22/20-1374.html>.
- [16] Brian Chmiel, Chaim Baskin, Ron Banner, **Evgenii Zheltonozhskii**, Yevgeny Yermolin, Alex Karbachevsky, Alex M. Bronstein, and Avi Mendelson. "Feature Map Transform Coding for Energy-Efficient CNN Inference". In: *International Joint Conference on Neural Networks (IJCNN)*. July 2020, pp. 1–9. DOI: 10.1109/IJCNN48605.2020.9206968. URL: <https://arxiv.org/abs/1905.10830>.
- [17] Yochai Zur, Chaim Baskin, **Evgenii Zheltonozhskii**, Brian Chmiel, Itay Evron, Alex M. Bronstein, and Avi Mendelson. "Towards Learning of Filter-Level Heterogeneous Compression of Convolutional Neural Networks". In: *ICML AutoML Workshop* (Apr. 2019). URL: <https://arxiv.org/abs/1904.09872>.
- [18] Chaim Baskin, **Evgenii Zheltonozhskii***, Tal Rozen, Natan Liss, Yoav Chai, Eli Schwartz, Raja Giryes, Alexander M. Bronstein, and Avi Mendelson. "NICE: Noise Injection and Clamping Estimation for Neural Network Quantization". In: *Mathematics* 9.17 (Sept. 2021): *Computational Optimizations for Machine Learning*. Ed. by Freddy Gabbay. ISSN: 2227-7390. DOI: 10.3390/math9172144. URL: <https://www.mdpi.com/2227-7390/9/17/2144>.
- [19] Chaim Baskin, Natan Liss, Eli Schwartz, **Evgenii Zheltonozhskii**, Raja Giryes, Alex M. Bronstein, and Avi Mendelson. "UNIQ: Uniform Noise Injection for Non-Uniform Quantization of Neural Networks". In: *ACM Transactions on Computer Systems* 37.1–4 (Mar. 2021). ISSN: 0734-2071. DOI: 10.1145/3444943. URL: <https://arxiv.org/abs/1804.10969>.
- [20] Chaim Baskin, Natan Liss, **Evgenii Zheltonozhskii**, Alex M. Bronstein, and Avi Mendelson. "Streaming Architecture for Large-Scale Quantized Neural Networks on an FPGA-Based Dataflow Platform". In: *IEEE International Parallel and Distributed Processing Symposium Workshops*. May 2018, pp. 162–169. DOI: 10.1109/IPDPSW.2018.00032. URL: <https://arxiv.org/abs/1708.00052>.