**Background.** Genome-wide association studies (GWAS) have identified thousands of variants linked to complex traits and diseases, yet distinguishing causal from non-causal variants in linkage disequilibrium (LD) remains challenging. Statistical fine-mapping can narrow candidate sets but often leaves ambiguity, especially in regions of high LD. GWAS are also underpowered to detect associations involving rare variants, and the mechanisms by which causal variants affect genome function in specific cell types are largely unknown. Integrating GWAS with functional genomics or evolutionary conservation can prioritize likely causal variants, but these approaches rely on broad genomic annotations that cannot separate causal from nearby non-causal variants.

Sequence-to-function (S2F) deep learning models address this limitation by predicting molecular measurements directly from DNA sequence at single-base resolution. By comparing predictions from reference and alternative alleles, S2F models can infer the molecular effects of variants, offering insight into both common and rare genetic variation. However, existing models are constrained to the experimental assays on which they were trained, which represents only a fraction of possible cellular contexts.

Here, we introduce DNACipher, a model that integrates long-sequence modeling with imputation capabilities to predict variant effects across a wide range of observed and unobserved biological contexts. We also present Deep Variant Impact Mapping (DVIM), a framework for prioritizing trait-associated variants based on DNACipher-predicted cell type-specific regulatory effects.

**Methods.** The DNACipher model takes three inputs: genomic sequence, one-hot encoded cell type, and one-hot encoded assay. Latent representations of cell type and assay are learned by the model, whereas the sequences are parsed through a previous pre-trained model (Enformer) to generate embeddings. These inputs are then passed into a multi-layer perceptron with ~1.3 million parameters. This design makes it possible to predict variant effects for all combinations of cell types and assays, including those which have not been observed. DVIM works by calling statistically significant common and rare variant DNACipher-effects at GWAS loci, by using non-trait associated common variants to generate a null distribution.

**Results.** After training on ENCODE data, DNACipher can predict functional variant effects in 38,582 distinct contexts (191 biological samples across 202 assays) over distances of 196 kb. We validate DNACipher on held-out experiments and fine-mapped GTEx expression quantitative trait loci (eQTLs), and demonstrate that it can detect >7-fold more molecular effects of high-confidence causal-eQTLs compared to alternative methods.

DNACipher DVIM applied to Type 1 Diabetes (T1D) GWAS reduced the average size of fine-mapping credible sets from 24 variants to 1.4 variants per GWAS signal, while simultaneously revealing cell type-specific effects that were corroborated by single-nucleus ATAC-seq data. We also tested DNACipher-DVIM variants using luciferase assays, and found variants predicted to be functional in T1D showed higher enhancer activity and greater allele-specific differences compared to predicted non-functional variants.

Furthermore, DNACipher DVIM predicted 6,547 rare variants as functional at T1D loci, all of which are previously uncharacterized in the context of the disease. Additional analyses showed these rare variants were significantly enriched for immune-related trait associations.

**Conclusion.** Overall, DNACipher-DVIM predicts common and rare causal variants and their precise functional effects at GWAS hit loci.