Cell Genomics

CellPress
OPEN ACCESS

Minireview

Causal effect estimation from *trans*-regulatory single-cell CRISPR screens

Oliver P. Christensen,[1,5] Alex Markham,[2,5] Hyunseung Kang,[3] Erin Gabriel,[4,5] and Tune H. Pers[1,5,*]
[1]Novo Nordisk Foundation Center for Basic Metabolic Research, University of Copenhagen, Copenhagen, Denmark
[2]Department of Mathematical Sciences, University of Copenhagen, Copenhagen, Denmark
[3]Department of Statistics, University of Wisconsin–Madison, Madison, WI, USA
[4]Section of Biostatistics, University of Copenhagen, Copenhagen, Denmark
[5]Pioneer Centre for SMARTbiomed, University of Copenhagen, Copenhagen, Denmark
*Correspondence: tune.pers@sund.ku.dk
https://doi.org/10.1016/j.xgen.2026.101251

Recent advances in single-cell transcriptomics and CRISPR-based genome editing have enabled large-scale perturbation experiments with genome-wide expression readouts. Single-cell CRISPR screens offer the opportunity to move beyond correlation and estimate causal effects of genetic perturbations on gene expression at scale. These approaches promise to substantially deepen insights into cellular functions and disease mechanisms. However, interpreting statistical associations as causal effects requires additional assumptions beyond those needed for standard statistical analyses. In this minireview, we introduce key concepts and principles for causal effect estimation in *trans*-regulatory single-cell CRISPR studies. We describe a set of assumptions under which estimates from existing statistical methods admit a causal interpretation and provide a concise overview of these approaches. Finally, through an illustrative example, we demonstrate how violations of these assumptions can bias estimated effects.

INTRODUCTION

Since the advent of single-cell sequencing, researchers have been able to estimate the causal effect of external interventions (e.g., drugs) on cell-type-resolved expression of a vast number of genes.[1,2] While approaches based on external (extracellular) interventions have greatly expanded our knowledge of cell biology and physiology, it is becoming increasingly important to establish causal relationships between perturbations of genes (e.g., knockdown of genes) and the expression levels of other genes in a given cell type. Traditionally, gene relationships have been approximated through analysis of gene co-expression patterns.[3] However, these relationships are defined by correlation rather than causation, and it can therefore be difficult to give them a causal interpretation.

Recently, thanks to rapid developments in genome editing techniques, pooled CRISPR screens combined with single-cell transcriptomic readouts (here referred to as single-cell CRISPR screens) offer a substantial step toward mapping causal relationships between perturbations and genes (Figures 1A and 1B).[4] Several flavors of single-cell CRISPR screening exist, including targeted Perturb-seq, direct capture Perturb-seq using CROP-seq, and related approaches, which are often collectively referred to as Perturb-seq.[4–6] Single-cell CRISPR screens can be divided into *cis*-regulatory and *trans*-regulatory. *Cis*-regulatory single-cell CRISPR screens target regulatory, non-coding parts of the genome to investigate *cis*-regulatory mechanisms.[7] The scope of this minireview is focused on causal effect estimation based on *trans*-regulatory single-cell CRISPR screen data.

In *trans*-regulatory CRISPR screens, a guide RNA (gRNA) pool is created, consisting of multiple gRNAs that each typically target the transcription start site of a gene. The resulting changes in gene expression enable investigation of *trans*-regulatory mechanisms governing gene regulation. Typically, several gRNAs are included for the same target gene. In low multiplicity of infection (MOI) versions of single-cell CRISPR screens, the gRNA pool is delivered to the cell culture or tissue with the aim of introducing a single gRNA per cell, resulting in a population of cells each carrying different gRNAs.[4] High MOI versions aim to increase the information gained for each cell and thereby increase the scalability of the experiment.[8] We refer to these as combinatorial single-cell CRISPR screens, but they are sometimes termed high-MOI single-cell CRISPR screens in the literature. This distinction is relevant in the analysis of single-cell CRISPR data and might influence what methods to use. In this minireview, we focus on single-guide single-cell CRISPR screens.

Upon gRNA pool construction, the gRNAs are delivered to cells expressing dCas9—a catalytically inactive Cas9 variant that can bind DNA without cutting it—which uses the delivered gRNA to locate the transcription start site of the target gene. The dCas9 is fused to effector proteins that can either increase (CRISPR activation/CRISPRa) or decrease (CRISPR interference/CRISPRi) gene expression.[4,5]

In CRISPRa experiments, the dCas9 is often fused with VP64, a strong synthetic transcriptional activation domain, which recruits general transcriptional machinery resulting in increased expression of the target gene. In CRISPRi experiments, the dCas9 is often fused to a KRAB repressor domain, which recruits chromatin remodeling proteins, resulting in a more closed

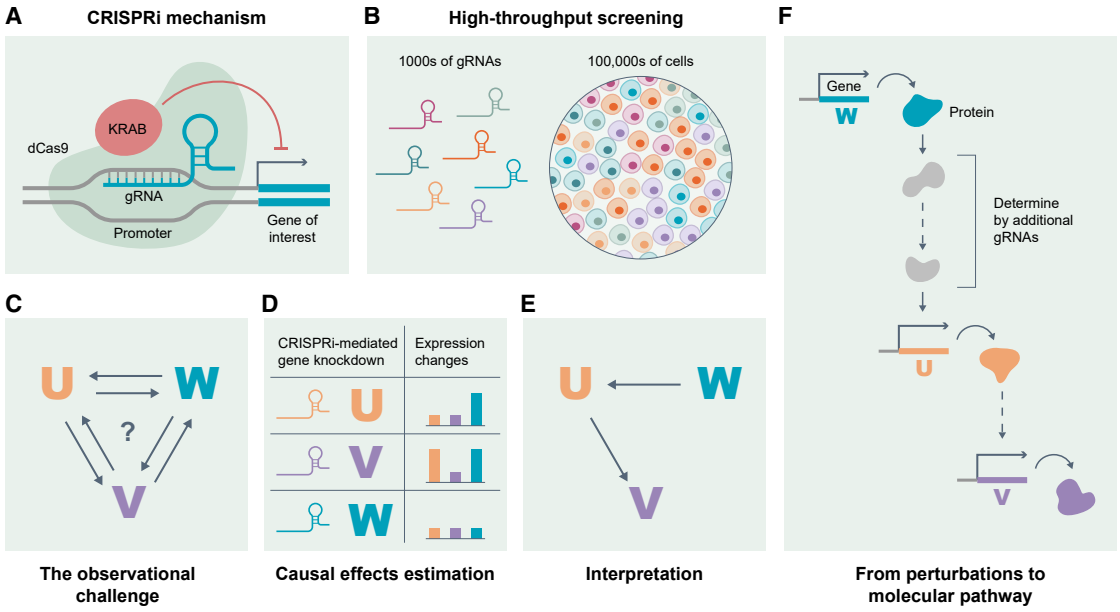

Figure 1. From CRISPR perturbations to causal interpretation and signaling pathway

(A) Schematic of the dCas9–KRAB CRISPR interference system: a catalytically dead Cas9 fused to the KRAB repressor is recruited by a gRNA to the promoter of a target gene, leading to repression of its transcription.

(B) Cartoon of a pooled gRNA library delivered into a cell population, where each cell ideally receives a single gRNA assigned at random and successfully binds to the desired transcriptional site.

(C) Observational association of three genes (U, V, and W) illustrating that, with only observational expression data, the underlying causal relationships are difficult to map.

(D) Illustration of how gRNAs targeting genes U, V, and W are used in single-cell CRISPR screens to knock down each gene (e.g., gene U) and subsequently measure expression changes of the others (e.g., genes V and W) to do causal effect estimation.

(E) Interpretation of the inferred causal effects knocking down genes U, V, and W. The ideal goal from interventional data is to infer causal relationships between U, V, and W whereby the causal effect between two genes can be inferred from the causal effect from knocking down a given gene on another measured gene. However, we can only identify the causal effect of the knockdown of one gene on another gene with the assumptions presented here and not the causal effect between two genes directly. See the Future perspectives section for further discussion.

(F) Conceptual illustration of how the estimated causal effects can be used to reconstruct signal transduction pathways at the protein level, linking genetic perturbations to downstream molecular pathway structure.

chromatin structure and thereby reduced expression of the target gene.[9,10] Finally, the cell culture or tissue is sequenced using a specialized single-cell transcriptomics protocol, which sequences not only the mRNA of each cell but also the gRNA present in a given cell.[4] The focus of this minireview is on CRISPRi experiments coupled with single-cell transcriptomics.

In statistical estimation, the association of variables is often considered a sufficient result of the analysis. This is not the case for causal effect estimation, where causal claims are the goal. Importantly, statistical methods do not yield causally interpretable results without additional causal assumptions, which need to be considered in the context of the experimental design and the subsequent analysis of the data. Notably, it is difficult to rely solely on observational measurements without introducing interventions that help prevent spurious correlations from being misinterpreted as causal effects (Figure 1C).[11] It is precisely in this way that single-cell CRISPR screens have the potential to enable researchers to map causal relationships between genes.

Existing reviews have provided broad overviews of single-cell CRISPR screen analysis methods[7] or focused on machine learning[12] and causal structure learning[13] applications. Although previous work has mentioned causal assumptions, specifically those in structure learning,[13] a more extensive review of the assumptions required for unbiased causal effect estimation is missing, especially assumptions that are not necessarily tied to a particular parametric model or mode of statistical analysis. Since valid causal assumptions are crucial for consistently estimating causal effects between perturbations and genes, we regard this as a critical gap in current practice.

Here, we first summarize fundamentals of causal reasoning in the context of *trans*-regulatory single-guide single-cell CRISPRi screens (henceforth "single-cell CRISPRi screens"). We then introduce commonly used causal effect estimation assumptions—consistency, no interference, conditional ignorability, and positivity—and discuss their relevance, as well as potential violations in single-cell CRISPRi screens. Next, we briefly describe the statistical characteristics of existing methods for analyzing such screens. Using a worked example, we illustrate the importance of the underlying data generation process satisfying a sufficient set of causal assumptions to achieve unbiased causal effect estimates. We conclude by providing perspectives for future work.

Box 1. Identifiability of the marginal causal rate ratio in *trans*-regulatory single-guide single-cell CRISPRi screens

Assuming consistency, no interference, and independent and identically distributed data, we define the model

$$(T, Z, Y(t)) \sim P,$$

where T is a binary treatment variable (control gRNA, gRNA targeting gene X), Z is the set of covariates (e.g., batch) sufficient for conditional ignorability, and $Y(t)$ is the potential outcome of the expression level of gene Y under treatment $T = t$. Here, P denotes the joint distribution over all observed and potential variables. We define a causal estimand relevant for *trans*-regulatory single-guide single-cell CRISPRi screens. In this minireview, we focus on the marginal causal rate ratio as defined below and show identifiability under law of iterated expectations (LIE), conditional ignorability (CI), positivity (PO), and consistency (CO):

$$\begin{aligned}\frac{\mathbb{E}(Y(1))}{\mathbb{E}(Y(0))} &= \frac{\mathbb{E}(\mathbb{E}(Y(1)|Z))}{\mathbb{E}(\mathbb{E}(Y(0)|Z))} && \text{(LIE)}\\ &= \frac{\mathbb{E}(\mathbb{E}(Y(1)|T=1,Z))}{\mathbb{E}(\mathbb{E}(Y(0)|T=0,Z))} && \text{(CI and PO)}\\ &= \frac{\mathbb{E}(\mathbb{E}(Y|T=1,Z))}{\mathbb{E}(\mathbb{E}(Y|T=0,Z))} && \text{(CO)}\end{aligned}$$

We arrived at the statistical estimand that is only a function of the observed data distribution. To further simplify the statistical estimand, we impose the following mean specification (MS) for the observed outcome:

$$\log(\mathrm{E}(Y|T=t, Z=z)) = \beta_0 + \beta_1 t + \beta_2 z.$$

Under (MS) and positivity, the statistical estimand simplifies to

$$\begin{aligned}\frac{\mathbb{E}(\mathbb{E}(Y|T=1,Z))}{\mathbb{E}(\mathbb{E}(Y|T=0,Z))} &= \frac{\mathbb{E}(e^{\beta_0+\beta_1+\beta_2 Z})}{\mathbb{E}(e^{\beta_0+\beta_2 Z})} && \text{(MS)}\\ &= \frac{e^{\beta_1}\,\mathbb{E}(e^{\beta_0+\beta_2 Z})}{\mathbb{E}(e^{\beta_0+\beta_2 Z})}\\ &= e^{\beta_1}.\end{aligned}$$

Thus, the marginal causal rate ratio equals e^{β_1}, which is estimated from a regression model.

FUNDAMENTALS OF CAUSAL EFFECT ESTIMATION IN SINGLE-CELL CRISPRi SCREENS

The central goal of causality is to reason about what would happen if one were to intervene on a system. In the simplest example, one might have two variables, T and Y, and then reason about what would happen to Y if one intervenes on T.

In causal effect estimation, we are interested in estimating the causal effect of T on Y. In general, the quantity we aim to estimate is called an estimand. A classical estimand is the average treatment effect. Under certain assumptions, this estimand is identifiable, meaning that we can express it using distributions we can observe. Specifically, the quantity can be estimated in perfect randomized experiments by comparing the average of Y of the treated samples to the average of Y of the control samples.[11,14–16] In the context of single-cell CRISPRi screens, this corresponds to investigating what would happen to the expression of gene Y if one were to introduce a gRNA targeting gene X in the cell. We will call this perturbation indicator T (Figure 1D).

Because transcript counts are often modeled using negative binomial regression with a log link or a log normal model with variance homogeneity assumptions, one model-based causal estimand of interest is the causal rate ratio. This corresponds to a fold change, which aligns with standard biological interpretation in differential expression analysis.[17] For this reason, we will focus on the causal rate ratio in what follows, although we acknowledge that other estimands do exist in the literature.[18]

The identifiability of a causal estimand is often considered independently of any specific parametric form, and the assumptions to associate the estimand with a particular estimator would only arise at the estimation stage if one were to choose a particular statistical model. Thus, the basic causal assumptions discussed herein (consistency, no interference, conditional ignorability, and positivity) are conceptually separate from any modeling choices made for statistical estimation.

The causal rate ratio is defined for the treatment of receiving a gRNA targeting gene X, without fixing the expression level of gene X to a specific value (Box 1). The estimand therefore does not represent the causal effect of gene X on gene Y. Instead, it quantifies the effect of the perturbation indicator T on gene Y. Although such a quantity may be related to the causal effect of gene X on gene Y, establishing this connection would require additional assumptions that are beyond the scope of this minireview. We discuss emerging opportunities and challenges in the Future perspectives section.

Single-cell CRISPR screen analysis goals

Single-cell CRISPRi screens typically serve one of two purposes, namely either discovery of novel cellular mechanisms or prediction of the effect of a given perturbation. Importantly, whether the ultimate goal is the former or the latter likely influences how causal relationships should be inferred in the initial analysis.[7]

Mechanistic understanding focuses on building explanatory models that deepen our knowledge of biological systems (Figure 1E). The results should be transparent to scientists and interpretability is therefore required. A concrete example of how causal relationships between perturbations and genes estimated from single-cell CRISPRi screen data can be useful is in mapping the components of a biological pathway and how these components regulate one another. Here, a pathway refers both to the regulators of a signal transduction pathway in the cell and to the gene program (i.e., the set of genes) regulated by that pathway.[19–21]

While signaling transduction predominantly happens at the protein level, single-cell CRISPRi screens intervene on the genome and measure the transcriptome. So, a leap from genes to protein has to be made to reason about pathways. Upon gene knockdown in *trans*-regulatory single-cell CRISPRi screens, the mRNA pool for the gene targeted by the gRNA is being depleted, which results in a reduction in the corresponding protein pool (Figure 1F). If knocking down a target gene leads to changes in the expression of many other genes, it is reasonable to infer that the target gene's protein is involved in gene regulation, except in cases where core components of the basal transcription machinery are targeted, which would be expected to have

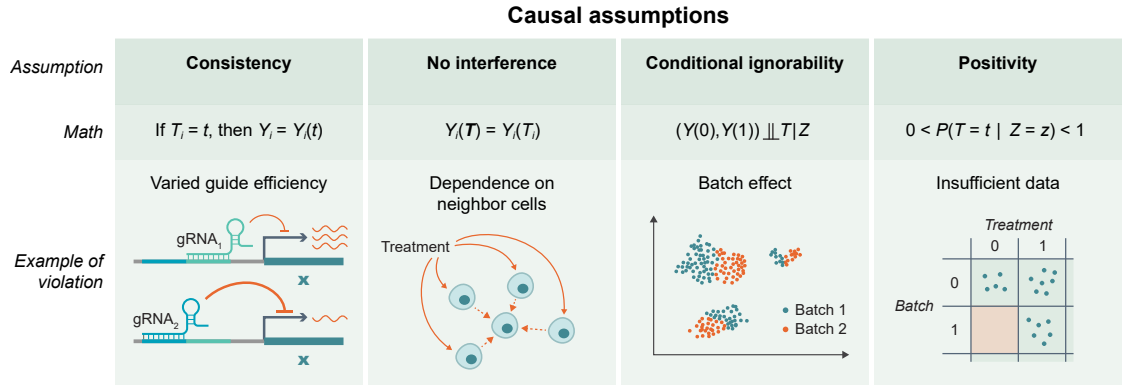

Figure 2. Schematic illustration of four causal assumptions allowing for causal interpretation of estimated association between perturbation and gene pairs in single-cell CRISPRi screens

For each assumption, the figure shows a simplified causal scenario in which the assumption is violated, leading to biased estimates of the causal effect of a gRNA targeting gene *X* on the expression of gene *Y*. Consistency violation is illustrated by aggregating gRNAs with different efficiencies into a single treatment indicator so that causal effect of the "treatment" (knockdown of gene *X*) corresponds to different underlying molecular interventions. The no interference violation is illustrated by a neighboring cell whose perturbation changes its state and, via paracrine or other intercellular signaling, alters the perturbation response of a nearby cell. Conditional ignorability violation is illustrated by a batch effect that could influence both the probability that a cell receives a given gRNA and the expression of gene *Y*, creating a spurious association that can be mistaken for a causal effect. Together, the panels highlight how each assumption connects to concrete features of single-cell CRISPRi screen experiments and how violations can compromise causal interpretations.

broad, non-specific effects on gene expression. If the knockdown of multiple target genes causes a similar change in the same gene set, then they are likely a part of the same gene regulation system. Consider cellular signaling transduction as an example of gene regulation where multiple proteins act together in a pathway to regulate specific sets of genes. In such pathways, a cascade of protein activation eventually leads to regulation of a transcription factor, which then ultimately leads to change in gene expression of downstream genes. In other words, perturbation effects are observed at the mRNA level and must be interpreted as indirect readouts of changes in protein activity and signaling state, rather than direct measurements of the signaling events themselves. Such mechanistic understanding of pathways and gene regulatory programs is not only a central goal of molecular research but also plays a crucial role in modern drug discovery.[7]

In contrast, perturbation prediction typically sets aside the requirement of interpretability and instead prioritizes predictive performance. Although such models may not initially provide an interpretable mechanistic understanding of gene regulation, they can still be valuable in biomedical research. This minireview focuses on causal effect estimation that can be used to gain mechanistic understanding. For reviews on predictive modeling, see, e.g., Tejada-Lapuerta et al.[13] and Gavriilidis et al.[12]

CAUSAL ASSUMPTIONS IN SINGLE-CELL CRISPRi SCREENS

Here, we introduce a set of basic common causal assumptions to infer causal relationships between perturbations and genes from single-cell CRISPRi screen data. We relate them to experimental design and data analysis considerations in single-cell CRISPRi screens and discuss strategies for handling potential violations of these assumptions (see Figure 2).

Besides causal assumptions, traditional statistical assumptions such as independent and identically distributed data and others are still necessary to ensure valid conclusions. These will not be discussed here, as the focus is on the causal assumptions underlying the move from statistical to causal interpretation.

Causal assumption 1—Consistency

Consistency requires that the treatment is well defined and has the same meaning across individuals.[11] For example, if we study the effect of a "10-mg dose of Drug A" on blood pressure, consistency assumes that "10 mg of Drug A" represents the same intervention for everyone: the same formulation, route of administration, and timing. If some people receive 10 mg orally once, while others receive 10 mg intravenously over several hours, then the "treatment" is not truly the same, and the causal effect of "10 mg of Drug A" is no longer well defined. In single-cell CRISPRi screens, a similar issue arises because gRNA efficiency can vary across different gRNAs targeting the same gene. As a result, the "perturbation of gene *X*" is not truly consistent if all gRNAs targeting that gene are simply seen as the same perturbation. One approach to mitigate this issue is to estimate causal effects at the gRNA level rather than aggregating at the target-gene level.[4]

A concrete example is using a single-cell CRISPRi screen to examine the signaling cascade downstream of the melanocortin 4 receptor (*MC4R*) gene in paraventricular hypothalamic MC4R neurons. Multiple gRNAs may target the *MC4R* gene, thereby reducing receptor expression and ultimately altering gene expression in the MC4R neuron. If the knockdown efficiency differs among the gRNAs, this could result in variable downstream changes in gene expression. In this case, the different gRNAs would have distinct causal effects, and treating them as the same intervention would violate the consistency assumption,

as the "treatment" would not be well defined with respect to the underlying biological perturbation.[22]

Causal assumption 2—No interference

The no-interference assumption states that one unit's treatment assignment should not affect another unit's outcome. A classic example is vaccination: if everyone around a person is vaccinated, their risk of infectious diseases decreases even if they themselves are not vaccinated, which violates the no-interference assumption.[16]

In single-cell CRISPRi screens, interference could occur if one cell (Cell 1) receives a gRNA that alters its state and subsequently modifies paracrine signaling, which in turn affects a nearby cell (Cell 2). In this case, the response of Cell 2 depends not only on the gRNA it receives but also on the gRNA received by Cell 1. This dependency violates the no-interference assumption, because the effect of a perturbation on one cell is influenced by the treatment of neighboring cells.

Hypothalamic MC4R neurons and pro-opiomelanocortin (POMC) neurons provide a concrete biological example of potential interference. POMC neurons signal to MC4R neurons via paracrine release α-MSH, which binds to the MC4R receptor and initiates a signal transduction cascade that ultimately leads to changes in gene expression. To examine the signaling transduction pathway downstream of MC4R, one would introduce a gRNA targeting *MC4R* and compare the resulting changes in gene expression using transcriptomic changes in MC4R neurons receiving the targeting versus control guide. Now consider a neighboring POMC neuron that receives a different gRNA reducing α-MSH levels. This would disrupt the paracrine signaling to the nearby MC4R neuron, preventing MC4R receptor activation, and inducing gene expression changes resembling those observed after directly perturbing the MC4R signaling cascade. In this scenario, the perturbation applied to the neighboring POMC neuron affects the outcome measured in the MC4R neuron—the causal effect of the gRNA in the MC4R neuron depends not only on its own perturbation but also on perturbations received by nearby cells, thereby violating the no-interference assumption.[22]

A potential solution is offered by optical pooled screens, where spatial information is preserved. Combined with spatial transcriptomics, this potentially allows modeling and adjusting for neighbor effects.[23–26] The consistency and no-interference assumptions are sometimes grouped into the stable unit treatment value assumption.

Causal assumption 3—Conditional ignorability

The conditional ignorability assumption requires that a sufficient set of confounders of T and Y are observed. This also implies that there should be no unmeasured confounding. Confounding leads to a non-causal association between T and Y, which can produce correlations that are mistaken for causal effects; for example, when disease status affects both treatment assignment and also the outcome.[14] A classic case is a kidney stone treatment study, where patients with small kidney stones were more likely to receive one treatment, while those with large stones were more likely to receive another.[27] Since stone size also strongly affects recovery rates, a naive comparison of treatments would suggest that the less effective treatment works better simply because it was given to patients with easier cases. Here, kidney stone size is the confounder that must be accounted for to correctly estimate the causal effect of treatment on recovery.[14]

In single-cell CRISPRi screens, the randomization of gRNAs across cells, in theory, should be sufficient to avoid confounding. This is exactly the property that makes single-cell CRISPRi screens a powerful experimental technique compared to traditional single-cell transcriptomics when the goal is to infer causal relationships between perturbations and genes. However, in practice confounding might still be present as illustrated in the next example.

Barry et al.[28] showed that batch effects confounded the causal effect of a non-targeting control gRNA—biasing its estimated effect on gene expression away from the true value of zero, despite the gRNA having no transcription start site target. This example illustrates the importance of adjusting for confounding in single-cell CRISPR screen data. Here, the batch effect acted as a confounder, but other types of confounding could also be present.

There are several ways of adjusting for confounding, the simplest being to include the confounder as a covariate in the model, which is possible for all the statistical association methods described later. However, the conditional ignorability assumption is fundamentally untestable, meaning that there is no way to prove the absence of unmeasured confounding. In practice, sensitivity analysis provides a partial remedy: by making assumptions about possible unmeasured confounding, one can test how robust causal effect estimates are to different levels of confounding. If results remain stable under strong hypothetical confounding, confidence in the estimates increases even though the assumption cannot be guaranteed.[29,30]

We stress that, although conditional ignorability is classically concerned with accounting for confounding by adjusting for specific pre-intervention variables to remove noncausal associations, conditioning on post-intervention variables that are affected by both gene Y and the gRNA targeting gene X will also violate conditional ignorability. This would result in biased causal effect estimation by introducing a non-causal association. Such post-intervention variables can be colliders; i.e., variables that—when conditioned on—open non-causal pathways. Additionally, even when post-intervention variables are not colliders, they may still be mediators of the causal effect, adjusting for which may not induce violations of ignorability but will still bias the estimate away from the targeted causal effect.

In the context of single-cell CRISPRi screens, this situation can arise if a gene Z is influenced by both gene Y and the gRNA targeting gene X. If gene Z, or some derivation of it, is included as a covariate in causal effect estimation, the conditional ignorability assumption is violated, and the resulting estimate of the causal effect of T on Y will be biased.

This is an important consideration, as many of the methods described later could inadvertently include post-intervention or outcome-derived covariates, such as gene expression levels of genes that may act as colliders.

Causal assumption 4—Positivity

Positivity requires that, for all combinations of treatments and confounders, there are individuals (or cells) with observed

Table 1. Overview of commonly used statistical methods for estimating perturbation-gene pair associations in single-cell CRISPR screen data

Method	Uni- vs. multivariate	Nonparametric inference	Outcome distribution	Parameter matrix rank	gRNA-efficiency aware	Statistical framework	Reference
SCEPTRE	univariate	yes	negative binomial	n/a	no	frequentist	Barry et al.[28]
Mixscale	univariate	no	negative binomial	n/a	yes	frequentist	Jiang et al.[19]
MAST	univariate	no	hurdle: logistic, Gaussian	n/a	no	frequentist	Finak et al.[31]
Normalisr	univariate	no	Gaussian	n/a	no	frequentist	Wang[32]
MIMOSCA	multivariate	yes	Gaussian	full rank	no	frequentist	Dixit et al.[4]
FR-perturb	multivariate	yes	Gaussian	low rank	no	frequentist	Yao et al.[8]
GSFA	multivariate	no	Gaussian	low rank	no	Bayesian	Zhou et al.[33]
scMAGeCK-LR	multivariate	yes	Gaussian	full rank	no	frequentist	Yang et al.[34]
Perturbation-response score	multivariate	yes	Gaussian	full rank	yes	frequentist	Song et al.[35]

Methods are classified as univariate or multivariate depending on whether they model one perturbation-gene pair at a time or several perturbation-gene pairs jointly. The "Nonparametric inference" column indicates whether uncertainty quantification (e.g., *p*-values or confidence intervals) is based on resampling or asymptotic assumptions. "Outcome distribution" specifies the assumed distributional form for gene expression measurements. If specified as Gaussian, then the method takes log-transformed data as input. "Parameter matrix rank" refers to the rank of the perturbation-effect parameter matrix, distinguishing full-rank models from low-rank factor-analytic approaches (specified as "n/a" for univariate methods). "gRNA-efficiency aware" indicates whether the method explicitly models heterogeneous gRNA efficiency on a single-cell level. "Statistical framework" indicates whether estimation and inference are carried out in a frequentist or Bayesian framework.

data. In single-cell CRISPRi screens, this means that there must be enough cells for each guide and within each level of confounders being adjusted for. Since the number of cells typically exceeds the number of guides, this condition is often satisfied.

OVERVIEW OF EXISTING ANALYSIS METHODS FOR SINGLE-CELL CRISPRi SCREENS

Several methods already exist for estimating statistical association between perturbations T and expression levels of gene Y. The common output of these methods is the fold change in gene expression of measured genes following the perturbation of a target gene. Many are differential expression methods tailored for single-cell CRISPR screens, each using distinct statistical strategies, briefly described below (see Table 1). Though not inherently causal, they can support causal inference—yet the required assumptions are rarely stated explicitly.

We note that to give a causal interpretation to the results of any of the methods described below, additional parametric modeling assumptions would be needed to allow for proper confounding adjustment. Additionally, we caution that for some methods, the standard implementations may automatically condition on post-intervention variables in the analysis or use them in some way in their pipeline, potentially inducing non-causal associations.

Univariate methods

Univariate methods analyze each perturbation-gene pair independently, modeling the effect of a single perturbation on one gene at a time. This approach simplifies the statistical modeling and inference, often enabling computationally efficient estimation. However, it inherently ignores dependencies between genes, limiting the ability to capture coordinated or indirect effects.

SCEPTRE

SCEPTRE[28] is a widely used method for analyzing single-cell CRISPR screen data. To estimate the gene expression change of a measured gene after a perturbation, it performs a univariate negative binomial regression for each pair of target and measured genes. A major advantage of SCEPTRE is its nonparametric inference of parameter estimates using permutation testing. This is implemented through a clever resampling algorithm that greatly reduces computational cost, making the method feasible on standard CPUs.

Mixscale

Mixscale[19] also uses a univariate negative binomial regression. A novel aspect of Mixscale is that it accounts for variability in perturbation effectiveness in each individual cell by first calculating a perturbation score for perturbed cells, which is then included as a covariate in the regression. This is a clear example of a post-intervention variable being included as a covariate, potentially inducing a non-causal association. The perturbation score is derived by projecting the gene expression of perturbed cells onto a vector representing the average difference between perturbed and control cells.

MAST

MAST[31] was originally developed for differential expression analysis in single-cell transcriptomics data but has also been applied to single-cell CRISPR screen data.[20] MAST performs univariate regression using a hurdle model that combines logistic and linear regression to account for zero inflation in single-cell data. It has been used in combination with consensus non-negative matrix factorization on single-cell CRISPR screen data, where consensus non-negative matrix factorization defines gene programs and MAST estimates the association between perturbations and these programs. When combined with matrix factorization, the regression approach resembles a multivariate regression.

Normalisr

Normalisr[32] was developed both to perform differential expression analysis in standard single-cell transcriptomics and to analyze single-cell CRISPR screen data. The novelty of the method lies primarily in the extensive normalization scheme proposed, which makes use of a digamma-based Bayesian log expression estimator. Nonlinear expansion of cell-level covariates is used to model mean effects. This can be another example of a post-intervention variable being included as a covariate, potentially inducing a non-causal association. Heteroskedastic variance modeling using similar nonlinear expansions is used to account for variance heterogeneity. Gene expression changes upon perturbations are then computed from the normalized data using Gaussian linear models, where Gaussian residual assumptions enable exact and computationally fast p-value calculations.

Multivariate methods

Multivariate methods jointly model multiple perturbation-gene pairs, allowing the estimation of effects while accounting for correlations and dependencies across genes. By leveraging the joint structure of the data, these approaches can improve estimation using regularization or low-rank assumptions and detect subtle perturbation effects that may be missed by univariate analysis.

MIMOSCA

MIMOSCA[4] was developed for one of the first single-cell CRISPR screen datasets. A key difference from the methods described above is that it performs multivariate regression, estimating all pairs of target and measured genes simultaneously. It applies elastic net regression on log-transformed transcript counts as a regularization strategy, thereby reducing the variance of the estimator and yielding more stable estimates. MIMOSCA also uses a nonparametric permutation testing strategy but leverages the complement set of gRNAs as controls due to the multivariate nature of the regression.

FR-perturb

FR-perturb[8] also applies multivariate regression on log-transformed transcript counts but assumes a low-rank structure in the parameter matrix. Specifically, sparse principal component analysis is first performed on the count matrix, followed by LASSO regression on one of the resulting matrices. This reduces the number of parameters to estimate, lowering estimator variance and improving efficiency given that the low-rank assumption holds. For statistical inference, FR-perturb relies on permutation testing.

GSFA

GSFA[33] also assumes a low-rank structure of the parameter matrix, combining factor analysis with multivariate regression, which is conceptually similar to FR-perturb. This reduces the number of estimated parameters. A key distinction is that GSFA uses Bayesian regression instead of frequentist regression. To achieve sparsity, it employs a spike-and-slab prior, and to further enforce sparsity in factor analysis, it applies a normal mixture prior.

scMAGeCK-LR

scMAGeCK-LR[34] is closely related to MIMOSCA, performing multivariate linear regression with regularization and permutation testing for statistical inference. However, it relies solely on ridge regression for regularization, rather than the elastic net (which combines LASSO and ridge regression).

Perturbation-response score

Perturbation-response score[35] is a method that extends scMAGeCK-LR to also estimate a cell-specific perturbation score. The method takes the perturbation effects estimated by scMAGeCK-LR and keeps them fixed while fitting individual perturbation scores for each cell-perturbation pair. Intuitively, the perturbation score can be interpreted as a cell-specific scaling of the mean perturbation effect of a gRNA. This scaling may be needed due to differences in guide efficiency in a specific cell and is shared across all measured genes. It should be noted, however, that the estimated mean perturbation effects are not changed and are therefore identical to those from scMAGeCK-LR. Perturbation-response score is therefore mainly relevant when variability in guide efficiency (or perturbation strength across cells) is of interest. We caution that depending on how perturbation response is derived, it may induce a spurious correlation.

Provided one does not induce a spurious correlation or condition on or use post-intervention variables in the above analyses, all of the above methods can provide estimates of some causal effect of the perturbation T on gene Y under the basic causal inference assumptions stated above. The additional specific assumptions depend on the specific analysis method and are beyond the scope of this minireview.

WORKED EXAMPLE OF CAUSAL ASSUMPTIONS

In this section, we provide scenarios of how violations of the causal assumptions from the previous section can bias the causal effect estimates in a hypothetical single-cell CRISPRi screen experiment. We show this for consistency, no interference, and conditional ignorability. We illustrate this using simple simulations of a pair of variables: a binary variable T indicating whether a cell has received a gRNA targeting gene X, and an arbitrary measured gene Y. For simplicity, the data are sampled from a Poisson distribution with a rate parameter determined by whether or not the gRNA targeting gene X is present. Rather than applying each method described earlier separately to the simulated data, we approximate their behavior using a simple Poisson regression that captures the core structure underlying many of the approaches described above. This is because most of these methods' distinguishing features concern how they handle realistic single-cell CRISPR screen data that deviate from a simple Poisson model, as well as other aspects not directly related to the causal assumptions considered here. In Poisson regression, the exponential of the estimated regression coefficient for the treatment can, under proper assumptions and correct modeling, be interpreted as the marginal causal rate ratio. In other words, the causal effect can be estimated as a fold change in gene expression.

In all simulations, the true causal effect of T on gene Y is set to a fold change of $\exp(-1)$, which corresponds to a roughly 63% decrease in expression (fold change ≈ 0.37). In all scenarios, the number of Monte Carlo replicates is set sufficiently high such that the variation from the Monte Carlo simulation is negligible. See https://github.com/perslab/Christensen-2026 for a script and documentation to rerun the examples.

Table 2. Estimated causal effects under different violations of causal assumptions in a simple simulation study

Assumption violation	Estimated causal effect (fold change)
None	0.368
Consistency	0.487
No interference	0.536
Conditional ignorability (confounder)	0.510
Conditional ignorability (collider)	0.490

We simulate a binary perturbation indicator for gene *X* and a gene expression outcome *Y* from a Poisson model, with the rate parameter determined by whether or not the gRNA targeting gene *X* is present. Across all scenarios, the true causal effect of perturbing *X* on *Y* is a fold change of $\exp(-1) \approx 0.368$, and causal effects are estimated using a Poisson regression model, interpreted as a fold change in mean expression. In all scenarios, the number of Monte Carlo replicates is set sufficiently high such that the variation from the Monte Carlo simulation is negligible. The row labeled "None" corresponds to data generated under a setting in which all three key assumptions (consistency, no interference, and conditional ignorability) are satisfied; in this case, the estimated causal effect closely matches the true value. The "Consistency" row represents a violation induced by aggregating gRNAs with different knockdown efficiencies into a single treatment indicator. The "No interference" row represents a violation where an environmental factor modulates the effect of the perturbation, analogous to neighboring cells altering the perturbation response of a knockdown. The "Conditional ignorability (confounder)" row corresponds to a violation caused by an unmeasured confounder that influences both the probability of receiving the perturbation and the expression level of *Y*. The "Conditional ignorability (collider)" row corresponds to a violation caused by conditioning on a collider that is affected by both the treatment assignment and the expression level of *Y*. For each scenario, the table reports the mean estimated causal effect, illustrating how even simple violations of the assumptions can induce substantial bias away from the true causal effect.

First, we simulated data in which all assumptions were satisfied. In these simulations, we were able to correctly estimate the true causal effect of perturbing gene *X* on the expression of gene *Y* (see Table 2). Then, we simulated data in which the consistency assumption was violated due to aggregating different gRNAs with differences in knockdown efficiency into the same treatment indicator so that "perturbing gene *X*" no longer represented the same intervention across cells. We next simulated data in which the no-interference assumption was violated due to an environmental variable affecting the causal effect, analogous to how the perturbation in neighboring cells could influence the perturbation response when knocking down a gene in a given cell. Then, we simulated data in which the conditional ignorability assumption was violated by introducing a common confounder that affects both the treatment assignment *T* and the expression of gene *Y*. Finally, we simulated data in which the conditional ignorability assumption was again violated, but this time by introducing a common collider that is affected by both the treatment assignment *T* and the expression of gene *Y*.

In all examples where a causal assumption was violated, we see a clear deviation from the true causal effect defined as a fold change of $\exp(-1) \approx 0.37$. Using biased causal effects as a basis for downstream analysis and interpretation could skew results and conclusions.

Taken together, these simulations illustrate how crucial it is that the true data-generating process satisfies the causal assumptions. When the goal is a causal interpretation of single-cell CRISPRi screen results, potential violations of these assumptions must therefore be carefully considered and addressed.

Demonstrating the consequences of violations of the underlying assumptions using real data is challenging, as the ground truth is not observed. One pragmatic approach is to leverage control guides and assume that the true causal effect on a given gene *Y* is zero. This strategy was employed by Barry et al.,[28] as described in the section on conditional ignorability.

FUTURE PERSPECTIVES

While many methods already exist for analyzing single-cell CRISPR screen data, there remains room to extend them to enable more rigorous causal interpretation. For instance, advanced techniques for confounder adjustment (such as inverse probability weighting,[36] augmented inverse probability weighting,[37] and targeted maximum likelihood estimation[38]) could be integrated into existing or new methods to complement the simpler approaches commonly used today (e.g., covariate adjustment). From an experimental-design perspective, it is crucial to measure, store, and properly adjust for as many covariates as feasible, including batch effects and other factors relevant to the specific experimental context, if one wishes to estimate a causal effect. However, it is equally important to have clear causal reasoning about all covariates or measures that are included or used in an analysis to avoid inducing non-causal associations. In many cases, existing or new methods could provide sensitivity bounds alongside estimated causal effects, allowing causal effects between perturbations and genes that appear more robust to be prioritized over others.

The consistency assumption also merits investigation in relation to single-cell CRISPRi screens, particularly regarding how to account for different types of treatment due to variation in gRNAs and their efficacy. Several methods attempt to address this,[4,19] but there is likely still room for improved approaches.

The no-interference assumption is another important concept for further investigation. It would be valuable both to quantify the degree of violation (e.g., the extent to which intercellular communication alters a cell's response to a gene knockdown) and to adapt existing methods that allow for such effects to single-cell CRISPR screen data analysis. This will likely require single-cell CRISPRi screens with spatial transcriptomics readouts to retain spatial information about cells; however, relevant experimental techniques are currently being developed.[25,26]

The scope of this minireview is limited to *trans*-regulatory single-guide single-cell CRISPRi screens. Assumption-driven identification of the marginal causal rate ratio could be extended to other causal estimands in related experimental setups, such as *cis*-regulatory screens, combinatorial perturbations, or other CRISPR-based techniques. In particular, combinatorial single-cell CRISPRi screens introduce the potential for interaction

effects for combinations of perturbations. These interaction effects need to be taken into account, because they can lead to bias in the estimation of the single-treatment effects.

Throughout, we have assumed that the treatment assignment of a given cell T is known. However, this is an idealization of reality: the gRNA is sequenced to determine the treatment a given cell received. This sequencing therefore produces a noisy estimate of the true treatment assignment. The concepts described here could be extended to account for this component of the true data-generating process.

Here, we have only discussed the causal effect of T on Y, when often the desired causal effect is of gene X on gene Y. It may be possible, under the assumptions of Mendelian randomization or instrumental variable analysis, to determine if there is an effect of gene X on gene Y, and under further parametric assumptions, to estimate that causal effect.[39,40]

In this paper, we have provided a very brief overview of existing statistical methods that, under certain sets of assumptions, can possibly be given a causal interpretation. However, this overview is by no means comprehensive, and we therefore refer readers to the original papers for detailed explanations of each method. A comprehensive comparison and benchmarking of the performance and limitations of these statistical methods within a causal context would be valuable future work.

Finally, several of the methods described here (e.g., FR-perturb,[8] GSFA[33]) primarily address the statistical power limitations often present in single-cell CRISPR screen studies. Such approaches are valuable because the experiments remain expensive, and obtaining sufficient data for traditional statistical analyses may be infeasible in many settings. Additional methods that build on these ideas will likely further improve the analysis of single-cell CRISPR screen data.

CONCLUSIONS

The field of CRISPR screens is currently strongly expanding, and diverse analysis strategies are still being explored. Here, we described the fundamental concepts of causal effect estimation in the context of single-cell CRISPRi screens. We discussed one set of common assumptions that allow for a causal interpretation of existing statistical methods: consistency, no interference, conditional ignorability, and positivity. Additionally, we emphasized that some set of assumptions are necessary to ensure a valid causal interpretation. Introducing each assumption with general examples and illustrations tailored to CRISPRi screens underlines that these are not abstract technicalities, but concrete conditions that must be scrutinized in practice.

To underscore the importance of the described causal assumptions, we presented a simulation-based worked example showing how causal effect estimates relying on these assumptions can become biased when any of these assumptions are violated. Our example illustrates that even seemingly reasonable analysis choices can lead to misleading conclusions when the underlying assumptions are not adequately justified.

Looking ahead, we envision opportunities for new experimental and computational methods that explicitly address potential violations of these key assumptions. Developing such approaches, and making their assumptions transparent, will be crucial for increasing the validity of causal interpretations of single-cell CRISPR screen analyses and for moving the field another step toward understanding the underlying causal systems of cellular biology.

ACKNOWLEDGMENTS

The Novo Nordisk Foundation Center for Basic Metabolic Research is an independent center at the University of Copenhagen, supported by an unrestricted donation from the Novo Nordisk Foundation (NNF23SA0084103 and NNF18CC0034900). O.P.C., A.M., E.G., and T.H.P. acknowledge the Pioneer Center for Statistical and Computational Methods for Advanced Research to Transform Biomedicine (SMARTbiomed), DNRF grant number P4.

A.M. was supported by Novo Nordisk Foundation grant NNF20OC0062897. T.H.P. received funding from the Lundbeck Foundation (R190-2014-3904), the Danish Council for Independent Research (grant number 8045-00091B), and the NIH (R01 DK124238). H.K. and E.G. would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program "Causal inference: From theory to practice and back again," where work on this paper was undertaken. This work was supported by EPSRC grant EP/Z000580/1.

DECLARATION OF INTERESTS

O.P.C. and T.H.P. are also affiliated with the Novo Nordisk Foundation Center for Genomic Mechanisms of Disease, Broad Institute of MIT and Harvard, Cambridge, MA 02142, USA. T.H.P. has received research support from Novo Nordisk A/S, his household holds shares in the company, and he has received counseling fees from Eli Lilly and Zealand Pharma A/S.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the authors used ChatGPT to improve the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

REFERENCES

1. Van de Sande, B., Lee, J.S., Mutasa-Gottgens, E., Naughton, B., Bacon, W., Manning, J., Wang, Y., Pollard, J., Mendez, M., Hill, J., et al. (2023). Applications of single-cell RNA sequencing in drug discovery and development. Nat. Rev. Drug Discov. *22*, 496–520, ISSN 1474-1784. https://doi.org/10.1038/s41573-023-00688-4.
2. Ludwig, M.Q., Cheng, W., Gordian, D., Lee, J., Paulsen, S.J., Hansen, S.N., Egerod, K.L., Barkholt, P., Rhodes, C.J., Secher, A., et al. (2021). A genetic map of the mouse dorsal vagal complex and its role in obesity. Nat. Metab. *3*, 530–545, ISSN 2522-5812. https://doi.org/10.1038/s42255-021-00363-1.
3. Crow, M., and Gillis, J. (2018). Co-expression in Single-Cell Analysis: Saving Grace or Original Sin? Trends Genet. *34*, 823–831, ISSN 0168-9525. https://doi.org/10.1016/j.tig.2018.07.007.
4. Dixit, A., Parnas, O., Li, B., Chen, J., Fulco, C.P., Jerby-Arnon, L., Marjanovic, N.D., Dionne, D., Burks, T., Raychowdhury, R., et al. (2016). Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. Cell *167*, 1853–1866.e17, ISSN 00928674. https://doi.org/10.1016/j.cell.2016.11.038.
5. Paul, D., Rendeiro, A.F., Schmidl, C., Krausgruber, T., Traxler, P., Klughammer, J., Schuster, L.C., Kuchler, A., Alpar, D., and Bock, C. (2017). Pooled CRISPR screening with single-cell transcriptome readout. Nat. Methods *14*, 297–301, ISSN 1548-7105. https://doi.org/10.1038/nmeth.4177.
6. Adamson, B., Norman, T.M., Jost, M., Cho, M.Y., Nuñez, J.K., Chen, Y., Villalta, J.E., Gilbert, L.A., Horlbeck, M.A., Hein, M.Y., et al. (2016). A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic

Dissection of the Unfolded Protein Response. Cell *167*, 1867–1882.e21, e21 ISSN 1097-4172. https://doi.org/10.1016/j.cell.2016.11.048.

7. Rood, J.E., Hupalowska, A., and Regev, A. (2024). Toward a foundation model of causal cell and tissue biology with a Perturbation Cell and Tissue Atlas. Cell *187*, 4520–4545, ISSN 0092-8674, 1097-4172. https://doi.org/10.1016/j.cell.2024.07.035.
8. Yao, D., Binan, L., Bezney, J., Simonton, B., Freedman, J., Frangieh, C.J., Dey, K., Geiger-Schuller, K., Eraslan, B., Gusev, A., et al. (2024). Scalable genetic screening for regulatory circuits using compressed Perturb-seq. Nat. Biotechnol. *42*, 1282–1295, ISSN 1087-0156. https://doi.org/10.1038/s41587-023-01964-9.
9. Qi, L.S., Larson, M.H., Gilbert, L.A., Doudna, J.A., Weissman, J.S., Arkin, A.P., and Lim, W.A. (2013). Repurposing CRISPR as an RNA-Guided Platform for Sequence-Specific Control of Gene Expression. Cell *152*, 1173–1183, ISSN 0092-8674. https://doi.org/10.1016/j.cell.2013.02.022.
10. Gilbert, L.A., Larson, M.H., Morsut, L., Liu, Z., Brar, G.A., Torres, S.E., Stern-Ginossar, N., Brandman, O., Whitehead, E.H., Doudna, J.A., et al. (2013). CRISPR-mediated modular RNAguided regulation of transcription in eukaryotes. Cell *154*, 442–451, ISSN 1097-4172. https://doi.org/10.1016/j.cell.2013.06.044.
11. Hernan, M.A., and Robins, J.M. (2020). Causal Inference: What If (Boca Raton: Chapman & Hall/CRC).
12. Gavriilidis, G.I., Vasileiou, V., Orfanou, A., Ishaque, N., and Psomopoulos, F. (2024). A mini-review on perturbation modelling across single-cell omic modalities. Comput. Struct. Biotechnol. J. *23*, 1886–1896, ISSN 20010370. https://doi.org/10.1016/j.csbj.2024.04.058.
13. Tejada-Lapuerta, A., Paul, B., Bauer, S., Aliee, H., Bengio, Y., and Theis, F.J. (2025). Causal machine learning for single-cell genomics. Nat. Genet. *57*, 797–808. https://doi.org/10.1038/s41588-025-02124-2.
14. Peters, J., Janzing, D., and Schölkopf, B. (2017). Elements of Causal Inference: Foundations and Learning Algorithms (Cambridge, MA, USA: MIT Press).
15. Judea, P. (2009). Causal inference in statistics: An overview. Stat. Surv. *3*, ISSN 1935-7516. https://doi.org/10.1214/09-SS057.
16. Imbens, G.W. (2024). Causal Inference in the Social Sciences. Annual Review of Statistics and Its Application *11*, 123–152, ISSN 2326-8298, 2326-831X. https://doi.org/10.1146/annurev-statistics-033121-114601.
17. Rosati, D., Palmieri, M., Brunelli, G., Morrione, A., Iannelli, F., Frullanti, E., and Giordano, A. (2024). Differential gene expression analysis pipelines and bioinformatic tools for the identification of specific biomarkers: A review. Comput. Struct. Biotechnol. J. *23*, 1154–1168, ISSN 20010370. https://doi.org/10.1016/j.csbj.2024.02.018.
18. Replogle, J.M., Saunders, R.A., Pogson, A.N., Hussmann, J.A., Lenail, A., Guna, A., Mascibroda, L., Wagner, E.J., Adelman, K., Lithwick-Yanai, G., et al. (2022). Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. Cell *185*, 2559–2575.e28, e28, ISSN 0092-8674, 1097-4172. https://doi.org/10.1016/j.cell.2022.05.013.
19. Jiang, L., Dalgarno, C., Papalexi, E., Mascio, I., Wessels, H.-H., Yun, H., Iremadze, N., Lithwick-Yanai, G., Lipson, D., and Satija, R. (2025). Systematic reconstruction of molecular pathway signatures using scalable single-cell perturbation screens. Nat. Cell Biol. *27*, 505–517. https://doi.org/10.1038/s41556-025-01622-z.
20. Schnitzler, G.R., Kang, H., Fang, S., Angom, R.S., Lee-Kim, V.S., Ma, X.R., Zhou, R., Zeng, T., Guo, K., Taylor, M.S., et al. (2024). Convergence of coronary artery disease genes onto endothelial cell programs. Nature *626*, 799–807, ISSN 1476-4687. https://doi.org/10.1038/s41586-024-07022-x.
21. Ota, M., Spence, J.P., Zeng, T., Dann, E., Milind, N., Marson, A., and Pritchard, J.K. (2025). Causal modelling of gene effects from regulators to programs to traits. Nature. https://doi.org/10.1038/s41586-025-09866-3.
22. Brüning, J.C., and Fenselau, H. (2023). Integrative neurocircuits that control metabolism and food intake. Science *381*, eabl7398, ISSN 0036-8075, 10959203. https://doi.org/10.1126/science.abl7398.
23. Ramezani, M., Weisbart, E., Bauman, J., Singh, A., Yong, J., Lozada, M., Way, G.P., Kavari, S.L., Diaz, C., Leardini, E., et al. (2025). A genome-wide atlas of human cell morphology. Nat. Methods *22*, 621–633, ISSN 1548-7105. https://doi.org/10.1038/s41592-024-02537-7.
24. Feldman, D., Singh, A., Schmid-Burgk, J.L., Carlson, R.J., Mezger, A., Garrity, A.J., Zhang, F., and Blainey, P.C. (2019). Optical Pooled Screens in Human Cells. Cell *179*, 787–799.e17, ISSN 00928674. https://doi.org/10.1016/j.cell.2019.09.016.
25. Binan, L.ïc, Jiang, A., Danquah, S.A., Valakh, V., Simonton, B., Bezney, J., Manguso, R.T., Yates, K.B., Nehme, R., et al. (2025). Simultaneous CRISPR screening and spatial transcriptomics reveal intracellular, intercellular, and functional transcriptional circuits. Cell *188*, 2141–2158.e18, e18. ISSN 0092-8674, 1097-4172. https://doi.org/10.1016/j.cell.2025.02.012.
26. Cheng, Y., Dang, S., Zhang, Y., Chen, Y., Yu, R., Liu, M., Jin, S., Han, A., Katz, S., and Wang, S. (2025). Sequencing-free whole-genome spatial transcriptomics at single-molecule resolution. Cell *188*, 6953–6970.e12, ISSN 00928674. https://doi.org/10.1016/j.cell.2025.09.006.
27. Charig, C.R., Webb, D.R., Payne, S.R., and Wickham, J.E. (1986). Comparison of treatment of renal calculi by open surgery, percutaneous nephrolithotomy, and extracorporeal shockwave lithotripsy. Br. Med. J. *292*, 879–882, ISSN 0267-0623. https://doi.org/10.1136/bmj.292.6524.879.
28. Barry, T., Mason, K., Roeder, K., and Katsevich, E. (2024). Robust differential expression testing for single-cell CRISPR screens at low multiplicity of infection. Genome Biol. *25*, 124, ISSN 1474-760X. https://doi.org/10.1186/s13059-024-03254-2.
29. Rosenbaum, P.R. (2005). Sensitivity Analysis in Observational Studies. In Encyclopedia of Statistics in Behavioral Science (John Wiley & Sons, Ltd). https://doi.org/10.1002/0470013192.bsa606.
30. VanderWeele, T.J., and Ding, P. (2017). Sensitivity Analysis in Observational Research: Introducing the E-Value. Ann. Intern. Med. *167*, 268–274, ISSN 0003-4819, 1539-3704. https://doi.org/10.7326/M16-2607.
31. Finak, G., McDavid, A., Yajima, M., Deng, J., Gersuk, V., Shalek, A.K., Slichter, C.K., Miller, H.W., McElrath, M.J., Prlic, M., et al. (2015). MAST: a flexible statistical framework for assessing transcriptional changes and characterizing heterogeneity in single-cell RNA sequencing data. Genome Biol. *16*, 278, ISSN 1474-760X. https://doi.org/10.1186/s13059-015-0844-5.
32. Wang, L. (2021). Single-cell normalization and association testing unifying CRISPR screen and gene co-expression analyses with Normalisr. Nat. Commun. *12*, 6395, ISSN 2041-1723. https://doi.org/10.1038/s41467-021-26682-1.
33. Zhou, Y., Luo, K., Liang, L., Chen, M., and He, X. (2023). A new Bayesian factor analysis method improves detection of genes and biological processes affected by perturbations in single-cell CRISPR screening. Nat. Methods *20*, 1693–1703, ISSN 1548-7091, 1548-7105. https://doi.org/10.1038/s41592-023-02017-4.
34. Yang, L., Zhu, Y., Yu, H., Cheng, X., Chen, S., Chu, Y., Huang, H., Zhang, J., and Li, W. (2020). scMAGeCK links genotypes with multiple phenotypes in single-cell CRISPR screens. Genome Biol. *21*, 19, ISSN 1474-760X. https://doi.org/10.1186/s13059-020-1928-4.
35. Song, B., Liu, D., Dai, W., McMyn, N.F., Wang, Q., Yang, D., Krejci, A., Vasilyev, A., Untermoser, N., Loregger, A., et al. (2025). Decoding heterogeneous single-cell perturbation responses. Nat. Cell Biol. *27*, 493–504, ISSN 1476-4679. https://doi.org/10.1038/s41556-025-01626-9.
36. Robins, J.M., Hernán, M.A., and Brumback, B. (2000). Marginal Structural Models and Causal Inference in Epidemiology. Epidemiology *11*, 550–560. ISSN 1044-3983.

37. Bang, H., and Robins, J.M. (2005). Doubly Robust Estimation in Missing Data and Causal Inference Models. Biometrics *61*, 962–973, ISSN 0006-341X. https://doi.org/10.1111/j.1541-0420.2005.00377.x.

38. Van Der Laan, M.J., and Rubin, D. (2006). Targeted Maximum Likelihood Learning. Int. J. Biostat. *2*, 1557–4679. https://doi.org/10.2202/15574679.1043.

39. Park, K., and Li, H. (2026). Confounder-robust causal discovery and inference in Perturb-seq using proxy and instrumental variables. arXiv:2601.01830.

40. Sun, Z., Kang, H., and Keles, S. (2026). Causal gene regulatory network inference from Perturb-seq via adaptive instrumental variable modeling. bioRxiv. https://doi.org/10.64898/2026.02.18.706642.