Large-scale composite hypothesis testing for omics analyses.

De Walsche Annaïg[1,2]✉, **Gauthier Franck**[2], **Boissot Nathalie**[3], **Charcosset Alain** [2], **and Mary-Huard Tristan** [1,2]

[1]MIA Paris-Saclay, INRAE, AgroParisTech, Université Paris-Saclay 91120 Palaiseau, France
[2]Génétique Quantitative et Evolution - Le Moulon, INRAE, CNRS, AgroParisTech, Université Paris-Saclay, 91190 Gif-sur-Yvette, France
[3]Génétique et Amélioration des Fruits et Légumes, INRAE, 84143 Montfavet, France

Abstract

Composite hypothesis testing using summary statistics is a well-established approach for assessing the effect of a single marker or gene across multiple traits or omics levels. Numerous procedures have been developed for this task and have been successfully applied to identify complex patterns of association between traits, conditions, or phenotypes. However, existing methods often struggle with scalability in large datasets or fail to account for dependencies between traits or omics levels, limiting their ability to control false positives effectively. To overcome these challenges, we present the qch_copula approach, which integrates mixture models with a copula function to capture dependencies between traits or omics, and provides rigorously defined p-values for any composite hypothesis. Through a comprehensive benchmark against eight state-of-the-art methods, we demonstrate that qch_copula controls Type I error rates effectively while enhancing the detection of joint association patterns. Compared to other mixture model-based approaches, our method notably reduces memory usage during the EM algorithm, allowing the analysis of up to 20 traits and $10^5 - 10^6$ markers. The effectiveness of qch_copula is further validated through two application cases in human and plant genetics. The method is available in the R package qch, accessible on CRAN.

meta-analysis | composite hypothesis testing | mixture model | summary statistics | pleiotropy | data integration

Correspondence: ***annaig.de-walsche@inrae.fr***

Introduction

Consider a study where the goal is to assess the joint effect of a specific drug treatment in two different tissues. One aims at defining a test procedure that rejects hypothesis H_0 «the drug has no joint effect», when both hypotheses H_0^1 «the drug has no effect on tissue 1» and H_0^2 «the drug has no effect on tissue 2» are false. This corresponds to a particular case of composite hypothesis testing where the composite H_0 hypothesis to be tested is $H_0^1 \cup H_0^2$. A popular strategy to perform Composite Hypothesis Testing (CHT) is to combine the test statistics and/or p-values derived for each of the marginal hypotheses H_0^1 and H_0^2 into a single summary statistics. While the p-values related to H_0^1 and H_0^2 can be obtained using common statistical procedures, providing a suitable summary statistic along with a valid rejection rule (that ensures the control of the false positive rate at the required nominal level) for the test of the composite H_0 hypothesis is not straightforward. The particular CHT problem of the form $H_0^1 \cup H_0^2$ was addressed as soon as the early 80s with the seminal work of (1), and its generalization to the case of testing $H_0^1 \cup \ldots \cup H_0^Q$ with $Q \geq 2$ has been investigated by (2).

In the context of omics studies CHT has rapidly become a standard procedure. In human genetics, it has been successfully used for e.g. mediation analysis (3–7), for the detection of genomic regions that are associated with two diseases such as Prostate Cancer and Type 2 Diabete (8). In these omics applications CHT is typically performed at the gene or marker level, resulting in a large number of simultaneously tested composite hypotheses. A first consequence is the need for CHT methods to explicitly provide p-values in order to be combined with multiple testing correction procedures. A second consequence is the opportunity to exploit the large amount of available data to model and infer the relationship between the 2 series of p-values, and to account for this relationship in the CHT procedure to better control Type I error rate. More recently, CHT has also been used in the context of integrative genomics to jointly analyse multiple omics traits, such as DNA methylation, copy number variation and gene expression (9, 10), or for detecting specific differential patterns in gene expression kinetics between two or more conditions (11). When applied to such studies, CHT typically requires the analysis of $Q > 2$ series of p-values (i.e. one serie per omics dataset) which led to the development of dedicated procedures.

Most existing methods rely on a similar model where the Q-uplet of p-values associated to each gene/marker is distributed as a multivariate mixture where each of the 2^Q components corresponds to a specific combination of H_0^q and H_1^q states (7, 9–11). While versatile, the mixture approach suffers from several limitations. First, as the number of components of the considered mixture model grows exponentially with Q, the number of series of p-values that can be handled by such CHT procedures is usually restricted to $Q \leq 10$ (9, 11) or less (e.g. $Q = 2$ and $Q = 3$ in (10) and (7), respectively). Second, while the type I error rate can be controlled through local or multi-class FDR in a mixture model framework (12–14), mixture model-based approaches do not provide p-values. This limitation hampers the application of most multiple testing procedures, restricts the use of diagnostic tools such as QQ-plots or histograms, and impedes the ability to compare with alternative methods. Lastly, only a few methods can efficiently handle both a number of p-value series Q greater than 2 and correlations between series, which may arise, for example, from using the same collection of samples to acquire the different omics measurements.

Based on the same mixture model approach, we propose a new method for CHT called qch_copula that explicitly accounts for the dependence structure across p-value series through a copula function (15). The new procedure comes with several important features. First we show how a consistent p-value can be defined from the mixture model approach and show the connection between adaptive Benjamini-Hochberg FDR control (16) applied to these p-values and a local FDR control (12, 13) directly applied to the posteriors obtained from the mixture model. Second, we provide a new implementation of the EM algorithm (17) that significantly alleviate the memory burden of the inference procedure of the approach, extending its application to series of p-values as large as $n = 10^5/10^6$ genes/markers and $Q = 20$. Third, we present the first extensive benchmark study of CHT procedures, comparing 8 recently developed procedures (5, 6, 8–11, 18). While half of the methods do not correctly control for Type I error rate when the p-value series are correlated, the qch_copula approach provides accurate Type I error rate control and yields excellent performance in terms of detection power. The procedure is then demonstrated through two application cases: the first in human genetics, where 14 association studies are jointly analyzed to identify new pleiotropic regions associated with psychiatric disorders, and the second in plant genetics, for detecting hotspot regions linked to resistance to multiple viruses in cucumber.

Materials and methods

This section introduces the framework of composite hypothesis testing then our new method for CHT based on a mixture model approach (Section Model). Inference and the testing procedure are presented in Section Inference and Section Testing composite hypothesis, respectively. The derivation of a memory-efficient EM algorithm is provided in Section Memory-Efficient EM algorithm. The simulation framework used to evaluate the performance of our procedure and its competitors are detailed in Section Simulation framework and Section Comparison with alternatives methods, respectively.

Composite hypothesis. Assume a collection of n items (e.g. genes or SNP) have been tested for their effects in Q conditions (e.g. traits, tissues or environments). We denote by H_0^q (resp. H_1^q) the null (resp. alternative) hypothesis corresponding to test q ($1 \leq q \leq Q$) and consider the set $\mathcal{C} := \{0,1\}^Q$ of all possible combinations of null and alternative hypotheses among Q. For a given configuration $c := (c_1, ..., c_Q) \in \mathcal{C}$, the joint hypothesis $\mathcal{H}^c$ is defined as

$$\mathcal{H}^c := \left(\bigcap_{q:c_q=0} H_0^q \right) \bigcap \left(\bigcap_{q:c_q=1} H_1^q \right)$$

Considering two complementary subsets $\mathcal{C}_0$ and $\mathcal{C}_1$ of $\mathcal{C}$, we define the composite null and alternative hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$ as

$$\mathcal{H}_0 := \bigcup_{c \in \mathcal{C}_0} \mathcal{H}^c, \quad \mathcal{H}_1 := \bigcup_{c \in \mathcal{C}_1} \mathcal{H}^c$$

We aim at testing $\mathcal{H}_0$ versus $\mathcal{H}_1$ for each of the n items.

Model. Let P_i^q denote the p-value obtained for test q on item i, and let $Z_i^q = -\Phi^{-1}(P_i^q)$ represent its negative probit transform, where Φ stands for the standard Gaussian cumulative distribution function (cdf). The vector $Z_i := (Z_i^1, ..., Z_i^Q)$ is referred to as the z-score profile of item i.
Each item i is associated with a latent vector $L_i := (L_i^1, ..., L_i^Q) \in \mathcal{C}$, where L_i^q is the binary variable indicating whether the null hypothesis H_{0i}^q ($L_i^q = 0$) or its alternative hypothesis H_{1i}^q ($L_i^q = 1$) holds. Under the assumption of independence between items, the z-score profile follows a mixture model with 2^Q components, expressed as:

$$Z_i \sim \sum_{c \in \mathcal{C}} w_c \psi^c. \tag{1}$$

where ψ^c is the distribution of Z_i conditional on $L_i = c$, and the $w_c's := \Pr\{L_i = c\}$ represent the mixing proportions.
Each component ψ^c corresponds to a multivariate distribution over $\mathbf{R}^Q$ that can be written in terms of univariate marginal distribution functions and a so-called copula function that describes the dependence structure between the Q z-scores (15):

$$\Psi_\theta^c(Z_i) = C_c(F_{c_1}^1(Z_i^1), ..., F_{c_q}^q(Z_i^q), ..., F_{c_Q}^Q(Z_i^Q))$$

where F_0^q (resp. F_1^q) is the marginal cdf of Z_i^q conditional on $L_i^q = 0$ (resp. $L_i^q = 1$). In what follows, we will assume that the copula function is common to all components and depends on a finite set of unknown parameters θ, that is, $C_c = C_\theta \ \forall c \in \mathcal{C}$. The corresponding density function for component c is then given by:

$$\psi_\theta^c(Z_i) = c_\theta \left(F_{c_1}^1(Z_i^1), ..., F_{c_Q}^Q(Z_i^Q) \right) \prod_{q:c_q=0} f_0^q(Z_i^q) \prod_{q:c_q=1} f_1^q(Z_i^q) \tag{2}$$

where f_0^q (resp. f_1^q) is the marginal density of Z_i^q conditional on $L_i^q = 0$ (resp. $L_i^q = 1$). Since z-scores are obtained from p-values whose distribution is known to be uniform over $[0,1]$ under H_0, all distributions are known: one has $F_0^q = \Phi$ for all q.
The expression obtained in Eq. (2) provides some hints about the complexity of the inference task: estimating the 2^Q emission distributions of the mixture model Eq. (1) actually reduces to determining the Q univariate cumulative distributions F_1^q and the copula parameter θ. In practice, one needs to choose a specific form of the copula distribution. Here we considered Gaussian copula due to its flexibility in specifying distinct correlation levels between each pair of variables. The density function of the Gaussian copula is given by:

$$c_\theta(u) = |\theta|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \xi_u^T (\theta^{-1} - I) \xi_u \right)$$

where $\xi_u = (\Phi^{-1}(u_1), ..., \Phi^{-1}(u_Q)) \quad \forall u \in [0,1]^Q$

and where θ is the $(Q \times Q)$ correlation matrix associated with the Gaussian copula.

Note that the scenario in which θ varies across components, and where the alternative distributions f_1^q are modeled as a Gaussian distribution, aligns with the methodology proposed by (10), which employs a mixture model of multivariate Gaussian distributions. Conversely, the specific case where $\theta = I_Q$ and where the alternative distributions f_1^q are inferred in a non-parametric way, corresponds to the independent mixture model introduced in (11).

Inference. The unknown parameters of model Eq. (1) are the distributions $f_1^1, ..., f_1^Q$, the copula parameter θ and the proportions w_c. Following (9) and (11), the inference procedure can be split into two steps:

1. Get an estimate of marginal densities f_1^q, $q = 1, ..., Q$;

2. Substitute the estimates $\hat{f}_1^q$ into Equation Eq. (1) and estimate both the proportions w_c and the copula parameter θ using maximum likelihood estimation.

Step 1: Inference of marginal distributions. Combining Model Eq. (1) with the definition of the ψ_c Eq. (2), one has

$$Z_i^q \sim \pi_q f_0^q + (1-\pi_q) f_1^q, \qquad \textbf{(3)}$$
$$\text{where} \qquad \pi_q = \sum_{c:c_q=0} w_c$$

that is the marginal distribution of Z_i^q is also a mixture model. Since $f_0^q = \phi$ for all q's, with ϕ the standard Gaussian density function, one needs to estimate f_1^q and π_q only.

The null proportions can be directly derived by applying the following estimator:

$$\hat{\pi}_q = [n(1-\lambda)]^{-1} |\{i : P_i^q > \lambda\}|,$$

where $\lambda \in [0,1]$ is a tuning parameter that can be determined through bootstrap (19). The estimated proportion $\hat{\pi}_q$ can then be plugged into equation Eq. (3), and the alternative distribution f_1^q can then be inferred through a nonparametric procedure using a kernel method (13). In terms of computational burden, this procedure can be executed for each of the Q components in parallel.

Step 2: Inference of the configuration proportions and the copula parameter. We now turn to the problem of inferring the copula parameter θ and the weights w_c using maximum likelihood estimation. Once the kernel estimates $\hat{f}_1^q$ are substituted in mixture model Eq. (1), the inference of the remaining parameters can be efficiently performed using a standard EM algorithm (17). In the present case, it is possible to obtain explicit update equations for both θ and the w_c's. Denoting $u_{ic} = (F_{c_1}(z_i^1), ..., F_{c_Q}(z_i^Q))^T$ and $\xi_{ic} = (\Phi^{-1}(u_{ic}^1), ..., \Phi^{-1}(u_{ic}^Q))^T$, one has :

$$\text{E step:} \qquad \hat{\tau}_{ic} = \widehat{\Pr}\{L_i = c | Z_i; \hat{\theta}\} = \frac{\hat{w}_c \psi_{\hat{\theta}}^c(Z_i)}{\sum_{c' \in \mathcal{C}} \hat{w}_{c'} \psi_{\hat{\theta}}^{c'}(Z_i)}$$

$$\text{M step:} \qquad \hat{w}_c = \frac{1}{n} \sum_i \hat{\tau}_{ic}$$

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n \sum_{c \in C} \hat{\tau}_{ic} (\xi_{ic} \xi_{ic}^T)$$

See Section 1 of the Supplementary Material for a detailed explanation of these derivations. Note that the expression of matrix θ ensures its positiveness. If applied directly, the (E) step involves computing and storing a posterior matrix of size $n \times 2^Q$. Although computational effort cannot be avoided, we demonstrate in Section Memory-Efficient EM algorithm that memory storage can be considerably reduced without loss of information.

Testing composite hypothesis. So far, methods relying on mixture model Eq. (1) (9–11) did not result in strictly valid test procedures as they did not provide a p-value. We present here how well-defined p-values can be derived in this framework.

Let c be the (unknown) configuration of the item under consideration, and $\mathcal{C}_0$ and $\mathcal{C}_1$ be any two complementary subsets of $\mathcal{C}$, and consider testing

$$\mathcal{H}_0 : \{c \in \mathcal{C}_0\} \quad \text{vs} \quad \mathcal{H}_1 : \{c \in \mathcal{C}_1\}$$

The mixture model Eq. (1) can be written in the following form:

$$Z \sim W_0 \psi^0 + W_1 \psi^1 := \psi$$

where $W_0 = \sum_{c \in \mathcal{C}_0} w_c$, $\psi^0 := W_0^{-1} \sum_{c \in \mathcal{C}_0} w_c \psi_c$ (resp. for W_1 and ψ^1). We consider the posterior

$$\begin{aligned} \tau(z) &= \Pr\{L \in \mathcal{C}_1 | Z = z\} \\ &= \sum_{c \in \mathcal{C}_1} \Pr\{L = c | Z = z\} \\ &= \frac{W_1 \psi^1(z)}{W_0 \psi^0(z) + W_1 \psi^1(z)} \end{aligned}$$

as a test statistic. The corresponding p-value is then:

$$\begin{aligned} \text{pval(z)} &= \Pr_{\mathcal{H}_0}\{\tau(Z^*) > \tau(z)\} \\ &= \int_{\mathbb{R}^Q} \mathbb{1}_{\{\tau(z^*) \geq \tau(z)\}} \psi^0(Z^*) dz^* \\ &= \frac{1}{W_0} \int_{\mathbb{R}^Q} \mathbb{1}_{\{\tau(z^*) \geq \tau(z)\}} (1 - \tau(z^*)) \psi(z^*) dz^* \end{aligned}$$

where the last equality comes from the definition of τ. In practice, estimates of posteriors

$$\hat{\tau}_i = \frac{\widehat{W}_1 \psi_{\hat{\theta}}^1(z_i)}{\widehat{W}_0 \psi_{\hat{\theta}}^0(z_i) + \widehat{W}_1 \psi_{\hat{\theta}}^1(z_i)}$$

can be computed for each item i. Using these estimates and approximating the integral by its empirical counterpart, one gets:

$$\widehat{\text{pval}}_i \quad = \frac{1}{n\widehat{W}_0} \sum_{j=1}^{n} \mathbb{1}_{\{\hat{\tau}_j > \hat{\tau}_i\}} (1 - \hat{\tau}_j).$$

Having access to p-values provides several advantages over existing methods that rely on posteriors only, from combining the testing procedure with any correction method for multiple testing to checking the p-value distribution for quality control and providing graphical displays such as Vulcano or Manhattan plots.

Beyond the empirical expression provided above, a theoretical equivalence can be established between p-values $\widehat{\text{pval}}_i$ corrected for multiple testing using the adaptive Benjamini Hochberg procedure (16) on one side, and the estimation of lFDR from the posteriors as presented in (12) and (13) on the other side. This equivalence is demonstrated in the Section 2 of the Supplementary Material.

Memory-Efficient EM algorithm. The classical EM algorithm (17) implementation requires the computation and the storage of the matrix $T = (\tau_{ic})$ in (E) step, which becomes cumbersome whenever n and/or Q are large: assuming e.g. $n = 10^5$ and $Q = 15$, the matrix T requires 26 GB of storage. This storage may be reduced by analyzing how these posteriors are used in the (M) step. For instance, the update of w_c at iteration $(t+1)$ can be reformulated as:

$$\hat{w}_c^{(t+1)} = \frac{1}{n} \sum_i \hat{\tau}_{ic} = \frac{1}{n} \sum_i \frac{\hat{w}_c^{(t)} \psi^c_{\hat{\theta}^{(t)}}(Z_i)}{S_i^{(t)}},$$

where $S_i^{(t)} = \sum_{c \in \mathcal{C}} \hat{w}_c^{(t)} \psi^c_{\hat{\theta}^{(t)}}(Z_i)$. Our reduced memory burden implementation of the EM algorithm works as follows: in the (E) step only the quantities $S_i^{(t)}, 1 \leq i \leq n$ are computed and saved; then the $\hat{w}_c^{(t)} \psi^c_{\hat{\theta}^{(t)}}(Z_i)$ are calculated on the fly for updating the weights $\hat{w}_c^{(t+1)}$ in the (M) step. A same procedure can be derived for the copula parameter $\hat{\theta}^{(t+1)}$ update, see Suppl. Mat. Section 3 for details.

This version of the EM algorithm retains the exact formulas of the original algorithm while significantly reducing its memory footprint from $O(n \times 2^Q)$ to $O(n + 2^Q)$. In our previous example ($n = 10^5$, $Q = 15$) the memory storage is downsized from 26 GB to 1 MB. This optimization makes it possible to apply our procedure to cases where n ranges from 10^5 to 10^6 and with up to $Q = 20$ conditions.

In what follows we will refer to the methodology presented so far (including mixture models and copulas, the use of an efficient EM algorithm for the inference and the derived p-values for the CHT procedure) as the qch_copula approach.

Simulation framework. We evaluated and compared the performance of our qch_copula method with existing methods in various simulated scenarios. Hereafter, we will refer to items and conditions as "markers" and "traits", mimicking a multi-trait genome-wide association study. Each simulated dataset consists of a $n \times Q$ matrix of z-scores generated as follows. Markers are clustered into $K+1$ clusters. Each cluster is characterized by a configuration c, i.e. all markers of cluster k share the same configuration $c_k = (c_{k1}, ..., c_{kQ})$. One has

$$Z_i \sim \mathcal{N}(2c_k, \Sigma_\rho),$$

where Σ_ρ is a $Q \times Q$ correlation matrix where all covariances between traits are set to ρ. Note that $\mu_{iq} = 0$ if H_0^q holds for item i, and $\mu_{iq} = 2$ otherwise. Assuming that matrix Σ_ρ is common to all clusters, a given cluster k is characterized by its associated number of markers and associated configuration vector. We considered different values for Q (2,8,16), two levels of correlation ($\rho = 0.3, 0.5$), and two scenarios: $sparse$ or $dense$. In the $sparse$ scenario, a large proportion (i.e. ≥ 0.9) of the markers were associated with none of the traits, whereas in the dense case, half of the markers were associated with at least one trait. For each setting (i.e., a combination of a number of traits, a correlation level and a scenario), simulations were repeated 25 times. In total, $3 \times 2 \times 2 = 12$ settings were investigated (see Table 6 in Suppl. Mat. for details). All simulations were performed using $n = 10^5$.

We provide the details for generating one dataset in Setting 1: $Q = 2$, $\rho = 0.3$, $sparse$, see Suppl. Mat. Section 4 for the description of other simulation settings. Twenty clusters of 50 markers associated with a single randomly chosen trait were simulated. Similarly, twenty clusters of 50 markers associated with the two traits were simulated. The process was repeated to generate ten clusters of 100 markers associated with one trait and ten clusters of 100 markers associated with two traits. A last cluster containing all remaining markers corresponding to the complete null configuration, i.e. $c = (0, 0)$, was generated. This last cluster included 96,000 markers, i.e., 96% of the total number of markers.

Comparison with alternative methods. Our benchmarking study evaluated eight recently proposed methods: DACT (5), HDMT (6), PLACO (8), adaFilter (18), IMIX (10), c-csmGmm (7), Primo (9) and qch (11). The first three methods - DACT, HDMT, and PLACO - use test statistics that combine two p-values per item and are restricted to scenarios involving no more than two test series. In contrast, IMIX, c-csmGMM, Primo and qch rely on mixture models similar to the one described in Equation Eq. (1) and can theoretically handle cases with more than two test series. However, in practice, both IMIX and c-csmGmm had prohibitive computational costs and/or the associated R package did not handle cases where $Q > 3$ and $Q > 2$,respectively. The adaFilter method takes a different approach, testing a specific composite null hypothesis of the form "fewer than q hypotheses among the Q tested are non-null", using an adaptive filtering multiple testing procedure with no restriction on Q. Among all considered methods, only PLACO, IMIX, Primo, c-csmGmm and qch_copula explicitly account for dependencies between p-values series. As Primo only provides an estimate of the local FDR for each marker, we relied on the

equivalence between local FDR (lFDR) estimation and p-values based on posteriors proven in Section Testing composite hypothesis to derive p-values for each marker for this method. Additionally, the data simulation process described in the previous section corresponds to the model described by Equations Eq. (1) and Eq. (2) with the marginal alternative distributions f_1^q being the Gaussian distribution $\mathcal{N}(2,1)$ and the copula parameter θ being Σ_ρ. This setup may be beneficial in terms of detection power to IMIX, Primo and qch_copula which rely on the same mixture model. However, all benchmarked methods should ensure a control of Type I error rate at the nominal level.

Results

Simulations. Our evaluation began by assessing the ability of the different methods to control Type I error (T1E) at the specified significance level and to detect true significant markers for $Q = 2$. We then extended the comparison to $Q = 8$ and $Q = 16$ for methods capable of handling larger cases.

Type I error rate control. In settings 1-4, $n = 10^5$ markers were tested for $Q = 2$ traits, with the objective of identifying markers associated with both traits. FDR control at a nominal level of 5% was applied using the adaptive BH procedure to generate the final list of candidate markers for each method. The results for the case $\rho = 0.3$ are summarized in Table 1, revealing significant departures from the nominal Type I error rate for most methods. As expected the qch_copula method exhibited substantial improvements over the original qch approach (which assumes conditional independence between p-value series), highlighting the importance of accounting for the dependency structures. More generally, only DACT_Efron, c-csmGmm and qch_copula consistently achieved proper FDR control, maintaining estimated FDR values close to 0.05 across all scenarios (settings 1-4). Similar results were observed when the correlation between traits was higher ($\rho = 0.5$; see Table 7 in Supp. Mat.), with increased FDR inflation among methods that failed to adequately control the FDR.

We further investigated TIE control in settings 5-8 and 9-12, corresponding to cases with $Q = 8$ and $Q = 16$, respectively, to identify markers associated with at least 2, 4 or 8 traits. Only Primo, adaFilter and qch_copula were included in these analyses, as the other methods are restricted to scenarios where $Q = 2$ or $Q \leq 3$. The results for $\rho = 0.3$ (settings 5-6 and 9-10) are presented in Table 2, while results for $\rho = 0.5$ (settings 7-8 and 11-12) can be found in Table 9 of the Supp. Mat.

In settings 5-8 ($Q = 8$), Primo, adaFilter, and qch_copula efficiently controlled for T1E, with both Primo and adaFilter showing a more conservative T1E compared to the target nominal level. In settings 9-12 ($Q = 16$), similar results were observed for adaFilter and qch_copula. A slight inflation of T1E was noted for qch_copula under the sparse scenario, although the actual risk typically stayed below 8% in most cases. We do not report results for Primo in settings 9-12 as it encountered significant computational challenges, either exhausting available memory or requiring excessively long runtimes, exceeding 24 hours.

Table 1. False discovery rate of the procedures for the sparse and dense scenarios with $Q = 2$ and $\rho = 0.3$ (settings 1 and 2).

Method	Sparse	Dense
DACT_Efron	0 (0)	0 (0)
DACT_JC	0 (0)	**0.08 (0.017)**
DACT_NO	0.032 (0.03)	0.014 (0.015)
HDMT	**0.255 (0.018)**	**0.129 (0.008)**
PLACO	**0.148 (0.091)**	0.058 (0.027)
adaFilter	**0.1 (0.015)**	**0.087 (0.008)**
qch	**0.256 (0.018)**	**0.145 (0.008)**
IMIX	**0.595 (0.128)**	**0.222 (0.094)**
c-csmGmm	0 (0)	0 (0)
Primo	0.046 (0.021)	**0.079 (0.01)**
qch_copula	0.045 (0.021)	0.057 (0.012)

Note: For each procedure the FDR (averaged over 25 runs) are displayed. Numbers in brackets correspond to standard errors. False discovery rate above 0.05 are in boldface.

Detection power. For simulations with $Q = 2$ and focusing on the methods that effectively controlled the FDR across all scenarios, DACT_Efron and c-csmGmm yielded highly conservative results, with detection power equal to zero. In contrast the qch_copula method achieved detection power ranging from 0.04 to 0.067 depending on the scenario, see Tables 7 and 8 in Suppl. Mat.

The results of the power detection analysis for settings 5-8 ($Q = 8$) and 9-12 ($Q = 16$) are presented in Table 2 for the three methods that were able to scale. In both the sparse and dense scenarios the performance of Primo was stable whatever the tested composite hypothesis $\mathcal{H}_0$, whereas adaFilter showed a marked decrease in power when moving from $\mathcal{H}_0$: "at least 2" to $\mathcal{H}_0$: "at least 8". This trend is consistent with the conservative behavior observed for adaFilter in the previous section. In contrast, qch_copula exhibited a substantial improvement in power over Primo and adaFilter regardless of the simulation setting, with e.g. a $\times 6$ increase of power compared to adaFilter in the sparse scenario ($Q = 16$, $\mathcal{H}_0$: "at least 8"). Similar results were observed under high correlation scenarios ($\rho = 0.5$, settings 7-8 and 11-12, see Table 9 in Suppl. Mat.).

In terms of computational performance, the average computational time for qch_copula with $Q = 16$ and $n = 10^5$ was 78 minutes for the model-fitting step and approximately 1 minute per tested composite hypothesis. The analysis was conducted on a computing platform with one thread and 3.225 GB of allocated RAM, running under the Ubuntu Linux operating system.

Application I: Detection of pleiotropic regions associated to psychiatric disorders. To illustrate our method, we performed a comprehensive analysis of 14 psychiatric disorders using data derived from genetic association studies conducted by the Psychiatric Genomics Consortium, see Section A and Table 10 in Suppl. Mat. for details. The

Table 2. Performance of the procedures Primo, adaFilter and qch_copula for the settings 5-6 ($Q = 8$) and 9-10 ($Q = 16$) with $\rho = 0.3$.

Q	Scenario	Test	Primo		adaFilter		qch_copula	
			FDR	Power	FDR	Power	FDR	Power
8	Dense	at least 2	0.008 (0.001)	0.143 (0.006)	0.018 (0.001)	0.317 (0.003)	0.051 (0.001)	0.609 (0.003)
		at least 4	0.005 (0.001)	0.126 (0.007)	0.006 (0.001)	0.148 (0.003)	0.044 (0.002)	0.534 (0.006)
		at least 8	0.048 (0.006)	0.125 (0.007)	0.017 (0.007)	0.033 (0.003)	0.05 (0.006)	0.186 (0.01)
	Sparse	at least 2	0.012 (0.004)	0.109 (0.008)	0.063 (0.005)	0.279 (0.008)	0.054 (0.004)	0.388 (0.007)
		at least 4	0.008 (0.004)	0.108 (0.009)	0.015 (0.004)	0.145 (0.005)	0.053 (0.004)	0.405 (0.014)
		at least 8	0.079 (0.019)	0.092 (0.009)	0.026 (0.021)	0.032 (0.01)	0.068 (0.023)	0.166 (0.019)
16	Dense	at least 2	-	-	0.02 (0.001)	0.278 (0.003)	0.052 (0.002)	0.646 (0.004)
		at least 4	-	-	0.008 (0.001)	0.165 (0.003)	0.057 (0.002)	0.678 (0.006)
		at least 8	-	-	0.006 (0.002)	0.094 (0.004)	0.057 (0.002)	0.659 (0.007)
	Sparse	at least 2	-	-	0.074 (0.006)	0.249 (0.005)	0.072 (0.004)	0.528 (0.008)
		at least 4	-	-	0.021 (0.005)	0.165 (0.007)	0.082 (0.004)	0.626 (0.01)
		at least 8	-	-	0.007 (0.005)	0.093 (0.006)	0.068 (0.006)	0.568 (0.018)

Note: For each procedure the FDR and the detection power (averaged over 25 runs) are displayed. Numbers in brackets correspond to standard errors.

initial data consists of p-values obtained for $n = 6,267,062$ to $12,438,502$ single nucleic polymorphisms (SNPs) for the different studies. The analysis aims to identify pleiotropic genomic regions simultaneously associated with multiple disorders. In the initial analysis of (20), the authors applied the MAGMA method (21) to aggregate p-values at the gene level before running their analysis, drastically reducing the size of the data to 26,024 genes. Aggregated data were analysed using PLACO (8) to detect genes associated with two or more disorders, focusing on the top genes detected in at least eight disorders. As the number of traits allowed by PLACO is limited to 2, all pairwise analyses of 2 among the 14 traits were performed. For a given gene, the list of associated disorders corresponded to those found significant in at least one pairwise analysis.

The initial approach suffers from several technical limitations. Firstly, while the aggregation step reduces the computational burden of the CHT procedure, it may lead to significantly less accurate outcomes. This is because some regions (e.g. intergenic regions or genes not listed in the annotation list) may not be represented, but also because the detection resolution is now limited to the gene level. Second, the use of a pairwise approach may yield ambiguous or inconsistent results. For example, gene TMX2 was detected in the ADHD-AN and the AN-SCZ analyses but not in the ADHD-SCZ analysis. Similarly, gene MIR2113, one of the two genes "that are identified in 10 psychiatric disorders", is declared significant in only 26 out of the 45 PLACO pairwise comparisons involving the ten candidate disorders, with Anorexia Nervosa (AN) being identified in only 1 out of its 13 associated pairwise comparisons. Finally, note that aggregating the results of all pairwise analyses does not come with any statistical guarantee regarding false positive control for the final gene list.

The goal of detecting pleiotropic SNPs exhibiting an effect in at least q disorders can be turned into a composite hypothesis testing procedure where the corresponding composite null hypothesis for SNP i is defined as follows:

$$\mathcal{H}_{0i} : \{\text{SNP } i \text{ is associated with at most } q-1 \text{ disorders}\}.$$

Table 3. Number of candidate genes and SNPs identified by PLACO and qch_copula, respectively, for different levels of pleiotropy (i.e., associated with at least the specified number of disorders).

At least # of disorders	8	9	10	11	>11
PLACO	38	21	2	0	0
qch_copula	1,608	498	211	25	0

Following the initial analysis, we focused on $q \geq 8$ and performed CHT at a nominal type-I error rate $\alpha = 0.05$. We ran qch_copula on the 5,172,884 SNPs common to all analyses and compared the outcomes with those obtained using the PLACO pairwise analyses. The results are displayed in Table 3 (see Suppl. Mat. Fig. 3 for quality control of the p-values distribution and Fig. S2 for the Manhattan plots).

Focusing on the composite hypothesis test with $q = 8$, our method identified 1,608 SNPs (see Fig 1 for their description), corresponding to 28 distinct regions. The regions identified by qch_copula may include more than one gene. Out of the 38 candidate genes detected using the PLACO approach, 35 were also found by our method (i.e. at least one significant SNP within the gene was identified). The three remaining genes NEGR1, TMX2, and C11orf31 that were identified exclusively by PLACO had only 7, 4, and 3 p-values below 0.01 out of 14 obtained from the MAGMA analysis (for each gene), respectively. Since the goal was to detect genes associated with a minimum of eight disorders, one would expect the number of significant MAGMA p-values to be close to eight or higher.

Additionally, qch_copula identified eight new regions, each including more than 10 SNPs, that were not detected by PLACO (see Table 4 for their description). Three of these eight regions were not detected in the initial analysis because the annotation list contained no gene located in these regions. Among them, the top region identified by qch_copula, located on chromosome 5 contained 338 SNPs (see the corresponding peak on the Manhattan plot in Fig 1, left). No-

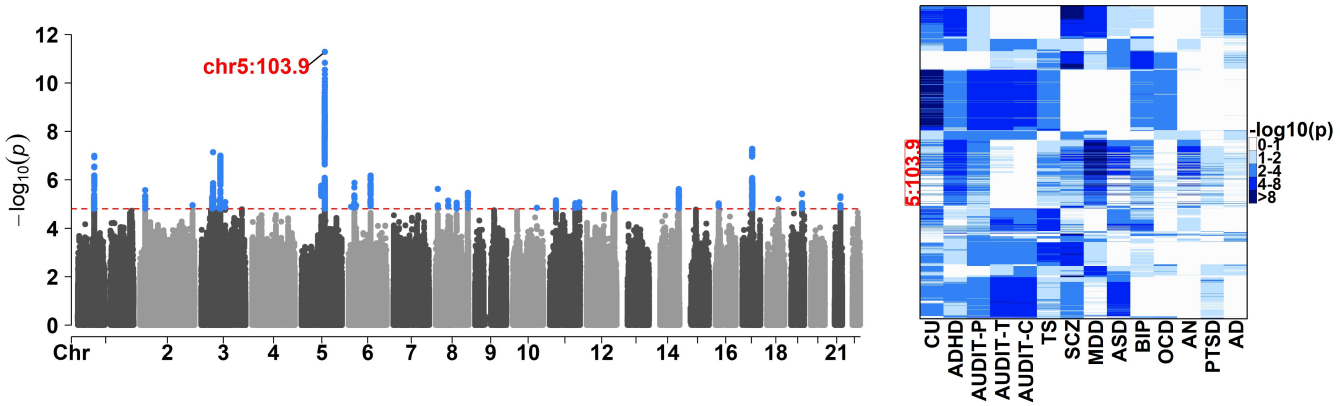

Fig. 1. Results of the composite hypothesis test identifying SNPs associated with at least eight disorders.
Left: -log10(p-values) of the composite hypothesis test along the chromosomes. The red dotted line represents the significance threshold at a nominal false discovery rate of 0.05. Significant SNPs are represented in blue. **Right:** Initial GWAS -log10(p-values) of the significant SNPs (in rows).

tably, 25 of these SNPs were detected as associated with 11 disorders (see the corresponding initial GWAS -log10(p-values) in Fig 1, right). This region shares an overlap with the gene RP11-6N13.1, which has been recurrently reported to be associated with psychiatric disorders (22–26), including ADHD, ASD, BIP, MDD, SCZ and TS, among others.

Table 4. Summary information of the eight new regions associated with at least eight disorders identified by qch_copula.

Chr	Position (Mb)	Nb SNPs	Top SNP	Pval top SNP
1	73.8 - 73.9	156	rs11210259	1.01e-07
2	**22.5 - 22.6**	14	rs11688810	2.66e-06
3	52.6 - 53.1	90	rs11922961	7.27e-08
5	**103.6 - 104.0**	338	rs30266	5.16e-12
6	25.9 - 26.3	12	rs115810	4.77e-06
8	143.3 - 143.3	11	rs4458978	3.41e-06
11	**27.6 - 27.6**	13	rs11030084	6.95e-06
11	113.2 -113.3	16	rs2734837	9.24e-06

Note: The regions not included in the initial analysis due to the absence of any gene in the annotation list are denoted in boldface.

Application II: Detection of viruses resistance hotspots regions in cucumber. Our second application case is based on GWAS summary statistics derived from the experiment described in (27). In this study, a panel of 226 cucumber elite lines, 40 landraces and 23 hybrids were inoculated with six viruses (denoted CGMMV, CMV, CVYV, PRSV, WMV and ZYMV hereafter) to evaluate their responses, see Section B in Suppl. Mat. for details. An individual GWAS was conducted on a number of SNPs ranging from $n = 378,049$ to $n = 424,393$ depending on the virus. The aim of the study was to identify QTLs associated with virus resistance in cucumber.

In the original analysis QTLs associated with resistance against multiple viruses were detected through a simple intersection of the lists of putative QTLs obtained from separate GWAS runs for each virus. As an alternative approach we addressed the detection of SNPs associated with resistance to at least q viruses through composite hypothesis testing. The corresponding null hypothesis for SNP i was defined as follows:

$$\mathcal{H}_{0i} : \{\text{SNP } i \text{ is associated with resistance at most } q-1 \text{ viruses}\}.$$

We focused on values $q \geq 2$ and applied the qch_copula method to the 339,804 SNPs common to all analyses, after excluding the genomic region spanning from 6.7 Mb to 12.9 Mb on chromosome 2, which corresponds to a non-recombinant region that produced anormalous p-value distributions. The nominal type-I error rate was set at $\alpha = 0.05$. Our analysis identified 1,845, 164, and 15 SNPs associated with resistance to at least two, three, and four viruses, respectively (see Supplementary Materials, Fig. 5 for quality control of p-value distributions, and Fig. 6 for Manhattan plots). For $q = 2$, the 1,845 SNPs mapped to five distinct genomic regions on chromosomes 1, 2, 5, and 6 (see Table 5 for details). All significant SNPs had at least two p-values below 10^{-4} out of the six GWAS (Fig 2). Analysis of the initial GWAS p-values indicates that the significant regions on chromosomes 1, 2, and 6 are associated with resistance to PRSV and ZYMV. In contrast, the region on chromosome 5 is associated with resistance to the remaining four viruses: WMV, CGMMV, CVYV and CMV. Notably, this region was also significant when testing for association with at least three or four viruses - no SNP being associated with resistance to more than four viruses.

Importantly, three of the detected hotspot regions were not reported in the original study (Table 5), of which two are strongly supported by external studies, confirming the hypothesis of a shared resistance mechanism across several viruses in these two regions. The hotspot region on chromosome 2 is associated with resistance to PRSV and ZYMV but also colocalizes with putative quantitative trait loci (QTLs) previously reported to confer resistance to CMV and CABYV (28). Similarly, the region on chromosome 6 (22.8–26.4 Mb) linked to resistance to PRSV and ZYMV, and colocalizes with QTLs associated with WMV (29) and CABYV (28). This highlights the enhanced statistical power of the joint analysis with a dedicated CHT procedure compared to individual GWAS approaches.

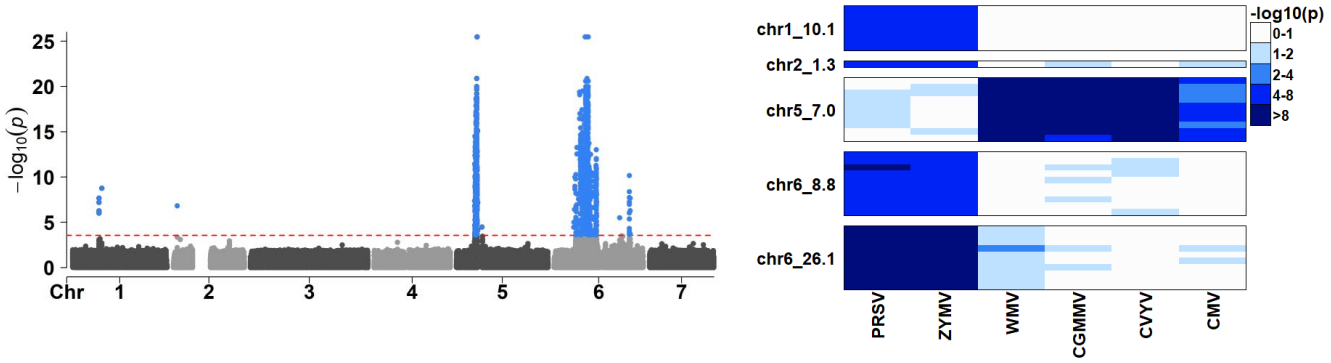

Fig. 2. Results of the composite hypothesis test identifying SNPs associated with resistance to at least two viruses.
Left: -log10(p-values) of the composite hypothesis test along the chromosomes. The red dotted line represents the significance threshold at a nominal false discovery rate of 0.05. Significant SNPs are represented in blue. **Right:** Initial GWAS -log10(p-values) of the top ten significant SNPs per regions identified (in rows).

Table 5. Summary of the five genomic regions associated with resistance to at least two viruses identified by qch_copula.

Chr	Position (Mb)	Nb SNPs	Top SNP	Pval top SNP
1	**9.1 - 10.1**	7	CucSaCL_Chr1_10190285	1.78e-9
2	**1.3 - 1.3**	1	CucSaCL_Chr2_01399444	1.54e-07
5	6.3 - 8.8	440	CucSaCL_Chr5_07197002	3.48e-26
6	6.8 - 14.7	1380	CucSaCL_Chr6_10814546	3.48e-26
6	**22.8 - 26.4**	16	CucSaCL_Chr6_26134362	6.81e-11

Note: Novel findings are highlighted in bold.

Conclusion

We have developed a novel approach called qch_copula for composite hypothesis testing based on a multivariate mixture model. Our method comes with a consistent p-value directly defined from the mixture model approach, which, to the best of our knowledge, is the first of its kind.

A key feature of qch_copula is its capacity to flexibly model the emission distributions, representing the probabilistic patterns by which each mixture component produces observable data. Previous published methods rely on fully parametric models, assuming both the f_0^q and f_1^q distributions of Equation Eq. (3) to be Gaussian density functions (7, 10), or (a mixture of) χ^2 density functions, after data transformation (9). While such modelings allow one to account for correlations across traits explicitly, they significantly constrain the shape of f_1^q. In comparison, qch_copula performs a non-parametric estimation of the alternative marginal distribution while still accounting for dependencies through the use of the copula function. This adaptive estimation procedure allows for a better fit of the data, resulting in an efficient control of Type I errors and improved detection performance. In practice, the correlation matrix is inferred jointly with the prior proportions within the (M) step of the EM algorithm, enhancing our ability to capture the dependence more accurately compared to methods like Primo and PLACO, where the correlation matrix is estimated upstream.

Our model assumes that the copula function is the same across all components, employing a single correlation matrix to model the dependence structure. While this may be seen as a limitation, it leads to practical benefits. Modeling the dependency component-wise, as in the IMIX procedure, can lead to a prohibitive computational burden and downgrade the results. This is because some components may be poorly represented in the dataset, compromising the inference of the associated correlation matrices. In contrast, qch_copula strikes a favorable balance between data fitting and computational efficiency in the inference procedure, offering a more practical and effective solution.

Our simulation study highlights the challenges in achieving a balance between Type I error control (T1E) and detection power in the presence of correlation across varying values of Q (2, 8, 16). Notably, most methods struggled to adequately control T1E under all scenarios - even when explicitly accounting for the dependence structures between the test series - while maintaining some detection power. In the $Q = 2$ case, qch_copula was the only method that both controlled FDR and maintained a non-zero detection power. For larger values of Q, qch_copula, Primo and adaFilter methods controlled FDR at the requested nominal level, but the last two methods exhibited low power levels for some testing hypotheses. Additionally, Primo suffered from prohibitively high computational costs when $Q = 16$. In comparison, qch_copula yielded substantial improvements in power across all simulation settings, even for larger values of Q (up to $Q \leq 20$).

Our procedure relies on summary statistics (i.e. the z-scores computed from the p-values of the initial test series). As such, it belongs to the larger family of meta-analysis procedures that have been developed in the last decade in the context of genetic association. However, composite hypothesis testing procedures differ from the many existing multi-trait methods based on summary statistics (30–36) that focus on testing the overall mean or variance of allelic effects over the traits. Indeed, CHT allows testing a wide range of composite hypotheses, as one can define any two complementary sets of configurations for $\mathcal{H}_0$ and $\mathcal{H}_1$. It is important to note that any series of composite hypotheses may be tested in parallel since the testing step stands independently of the model inference step and does not require parameter re-estimation. Importantly, relying on series of p-values exclusively makes our procedure amenable to data integration of studies with diverse outcomes, including continuous, binary, and time-to-event measurements. The exclusive reliance on p-values is also particularly advantageous when direct access to the raw

data is constrained due to ethical considerations or data confidentiality.

The proposed methodology is implemented in the R package `qch`, which is publicly available on CRAN.

ACKNOWLEDGEMENTS

The authors would like to thank the Psychiatric Genomics Consortium (PGC) for making the summary statistics of psychiatric disorders GWAS publicly available. Additionally, the authors sincerely thank Séverine Monnot for sharing the summary statistics of the cucumber virus GWAS and for her valuable assistance in interpreting the results of our analysis. The authors are grateful to the INRAE MIGALE bioinformatics facility (MIGALE, INRAE, 2020. Migale bioinformatics Facility, doi: 10.15454/1.5572390655343293E12) for providing computing resources.

FUNDING

This work was partially supported by INRAE's metaprogram DIGIT-BI0 and by KWS. The funders had no role in the conceptualization, design, data collection, analysis, publication decision, or manuscript preparation.

No competing interest is declared.

DATA AVAILABILITY

The psychiatric disorders GWAS summary statistics are available in the Psychiatric Genomics Consortium repository at https://pgc.unc.edu/for-researchers/download-results/, and can be accessed with the identifiers: an2017, anx2016, adhd2019, asd2019, sud2019-alcuse, bip2019, mdd2018, ocd2018, ptsd2019, scz2019asi and ts2019. The cucumber viruses resistance GWAS summary statistics are available in the following data repository: https://doi.org/10.57745/XQ3P72. Functions to perform composite hypothesis procedure can be found in the R package qch, available at https://cran.r-project.org/web/packages/qch/index.html.

Bibliography

1. M. E. Sobel. Asymptotic confidence intervals for indirect effects in structural equation models. *Source: Sociological Methodology*, 1982.
2. R. L. Berger. Multiparameter hypothesis testing and acceptance sampling. *Technometrics*, 1982.
3. Y.-T. Huang. Joint significance tests for mediation effects of socioeconomic adversity on adiposity via epigenetics. 2018.
4. Y.-T. Huang. Genome-wide analyses of sparse mediation effects under composite null hypotheses. 2019.
5. Z. Liu et al. Large-scale hypothesis testing for causal mediation effects with applications in genome-wide epigenetic studies. *Journal of the American Statistical Association*, 2021.
6. J. Y. Dai, J. L. Stanford, and M. LeBlanc. A multiple-testing procedure for high-dimensional mediation hypotheses. *Journal of the American Statistical Association*, 2022.
7. R. Sun, Z. R. McCaw, and X. Lin. Testing a large number of composite null hypotheses using conditionally symmetric multidimensional gaussian mixtures in genome-wide studies. *Journal of the American Statistical Association*, pages 1–22, 11 2024.
8. D. Ray and N. Chatterjee. A powerful method for pleiotropic analysis under composite null hypothesis identifies novel shared loci between type 2 diabetes and prostate cancer. *PLoS Genetics*, 2020.
9. K. J. Gleason et al. Primo: Integration of multiple gwas and omics qtl summary statistics for elucidation of molecular mechanisms of trait-associated snps and detection of pleiotropy in complex traits. *Genome Biology*, 2020.
10. Z. Wang and P. Wei. Imix: a multivariate mixture model approach to association analysis through multi-omics data integration. *Bioinformatics*, 2021.
11. T. Mary-Huard, S. Das, I. Mukhopadhyay, and S. Robin. Querying multiple sets of p-values through composed hypothesis testing. *Bioinformatics*, 12 2021.
12. G. J. McLachlan, R. W. Bean, and L. B. T. Jones. A simple implementation of a normal mixture approach to differential gene expression in multiclass microarrays. *Bioinformatics*, 2006.
13. S. Robin, A. Bar-Hen, J. J. Daudin, and L. Pierre. A semi-parametric approach for mixture models: Application to local false discovery rate estimation. *Computational Statistics & Data Analysis*, 2007.
14. T. Mary-Huard, V. Perduca, M.-L. Martin-Magniette, and G. Blanchard. Error rate control for classification rules in multiclass mixture models. *The international journal of biostatistics*, (2), 2022.
15. A. Sklar. Random variables, joint distribution functions, and copulas. *Kybernetika*, (6), 1973.
16. Y. Benjamini and Y. Hochberg. On the adaptive control of the false discovery rate in multiple testing with independent statistics. *Journal of Educational and Behavioral Statistics*, 2000.
17. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1): 1–22, 1977.
18. J. Wang et al. Detecting multiple replicating signals using adaptive filtering procedures. *The Annals of Statistics*, 50(4):1890 – 1909, 2022.
19. J. D. Storey, J. E. Taylor, and D. Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *J. R. Statist. Soc. B*, 2004.
20. H. Lu et al. A comprehensive gene-centric pleiotropic association analysis for 14 psychiatric disorders with gwas summary statistics. *BMC Medicine*, 2021.
21. C. A. de Leeuw, J. M. Mooij, T. Heskes, and D. Posthuma. Magma: Generalized gene-set analysis of gwas data. *PLoS Computational Biology*, 2015.
22. P. H. Lee et al. Genomic relationships, novel loci, and pleiotropic mechanisms across eight psychiatric disorders. *Cell*, (7), 2019.
23. L. Xia, K. Xia, D. R. Weinberger, and F. Zhang. Common genetic variants shared among five major psychiatric disorders: a large-scale genome-wide combined analysis. *Glob Clin Transl Res*, 2019.
24. Y. S. Lin, C. C. Wang, and C. Y. Chen. Gwas meta-analysis reveals shared genes and biological pathways between major depressive disorder and insomnia. *Genes*, 10 2021.
25. J. Ward et al. The genomic basis of mood instability: identification of 46 loci in 363,705 uk biobank participants, genetic correlation with psychiatric disorders, and association with gene expression and function. *Molecular Psychiatry*, 11 2020.
26. V. Powell et al. Investigating regions of shared genetic variation in attention deficit/hyperactivity disorder and major depressive disorder: a gwas meta-analysis. *Scientific Reports*, 12 2021.
27. S. Monnot et al. Unravelling cucumber resistance to several viruses via genome-wide association studies highlighted resistance hotspots and new qtls. *Horticulture Research*, 9, 2022.
28. S. Monnot et al. Genome-wide association studies to assess genetic factors controlling cucumber resistance to cabyv and cmv in crop fields and the attractiveness for their aphis gossypii vector. *Horticulture Research*, 2025.
29. Y. Wang et al. Molecularly tagged genes and quantitative trait loci in cucumber with recommendations for qtl nomenclature. *Horticulture Research*, 7:3, 2020.
30. A. Cichonska et al. Metacca: Summary statistics-based multivariate meta-analysis of genome-wide association studies using canonical correlation analysis. *Bioinformatics*, 7 2016.
31. Z. Wang, Q. Sha, and S. Zhang. Joint analysis of multiple traits using "optimal" maximum heritability test. *PLOS one*, 2016.
32. G. Qi and N. C. Id. Heritability informed power optimization (hipo) leads to enhanced detection of genetic associations across multiple traits. *PLOS Genetetics*, 2018.
33. D. Ray and M. Boehnke. Methods for meta-analysis of multiple traits using gwas summary statistics. *Genetic Epidemiology*, 3 2018.
34. P. Turley et al. Multi-trait analysis of genome-wide association summary statistics using mtag. *Nature genetics*, 2 2018.
35. H. Julienne et al. Jass: command line and web interface for the joint analysis of gwas results. *NAR Genomics and Bioinformatics*, 3 2020.
36. A. De Walsche et al. metage: Investigating genotype x environment interactions through gwas meta-analysis. *PLOS Genetics*, 21:1–26, 01 2025.
37. H. Hult and F. Lindskog. Multivariate extremes, aggregation and dependence in elliptical distributions. 2002.
38. L. Duncan et al. Significant locus and metabolic genetic correlations revealed in genome-wide association study of anorexia nervosa. *American Journal of Psychiatry*, 2017.
39. T. Otowa et al. Meta-analysis of genome-wide association studies of anxiety disorders. *Molecular psychiatry*, 10 2016.
40. D. Demontis et al. Discovery of the first genome-wide significant risk loci for attention-deficit/hyperactivity disorder. *Nature genetics*, 2019.
41. J. Grove et al. Identification of common genetic risk variants for autism spectrum disorder. *Nature genetics*, 2019.
42. S. Sanchez-Roige et al. Genome-wide association study meta-analysis of the alcohol use disorders identification test (audit) in two population-based cohorts. *American Journal of Psychiatry*, 2019.
43. E. A. Stahl et al. Genome-wide association study identifies 30 loci associated with bipolar disorder. *Nature genetics*, 2019.
44. J. A. Pasman et al. Gwas of lifetime cannabis use reveals new risk loci, genetic overlap with psychiatric traits, and a causal influence of schizophrenia. *Nature neuroscience*, 2018.
45. N. R. Wray et al. Genome-wide association analyses identify 44 risk variants and refine the genetic architecture of major depression. *Nature genetics*, 2018.
46. P. D. Arnold et al. Revealing the complex genetic architecture of obsessive-compulsive disorder using meta-analysis. *Molecular Psychiatry*, 2018.
47. C. M. Nievergelt et al. International meta-analysis of ptsd genome-wide association studies identifies sex- and ancestry-specific genetic risk loci. *Nature Communications 2019 10:1*, 2019.
48. M. Lam et al. Comparative genetic architectures of schizophrenia in east asian and european populations. *Nature genetics*, 2019.
49. D. Yu et al. Interrogating the genetic determinants of tourette's syndrome and other tic disorders through genome-wide association studies. *American Journal of Psychiatry*, 2019.

Supplementary Note 1: Estimation of the proportions w_c and the copula parameter θ using an EM algorithm.

A. Initialisation. The prior proportions are initialized using

$$\hat{w}_c = \prod_{q:c_q=0} \hat{\pi}_0^q \prod_{q:c_q=1} (1-\hat{\pi}_0^q)$$

where $\hat{\pi}_0^q$ corresponds to the marginal prior probability in model (3) estimated in the first step of the inference procedure. Parameter θ is initialized using the inversion of Kendall's tau estimator (37):

$$\hat{\theta}_{ij} = \sin \frac{\pi}{2} \tau_{ij}$$

where τ_{ij} is the Kendall's tau for the random vector $(Z_i, Z_j)^T$.

B. Maximum likelihood estimate for the Gaussian copula parameter. The complete likelihood of the observations $z_i = (z_i^1, ..., z_i^Q), 1 \leq i \leq n$ and the latent variables $l_{ic} = \mathbb{1}_{\{l_i=c\}}, 1 \leq i \leq n, c \in \mathcal{C}$ is given by:

$$\begin{aligned} L(w,\theta|z,l) &= \prod_{i=1}^n \prod_{c\in\mathcal{C}} [w_c \psi_\theta^c(z_i)]^{l_{ic}} \\ &= \prod_{i=1}^n \prod_{c\in\mathcal{C}} \left[w_c \left(\prod_{q:c_q=0} \phi(z_i^q) \prod_{q:c_q=1} f_1^q(z_i^q) \right) c_\theta(F_{c_1}(z_i^1), ..., F_{c_Q}(z_i^Q))) \right]^{l_{ic}} \end{aligned}$$

Define $Q\left(w,\theta|w^{(t)},\theta^{(t)}\right)$ as the expected value of the complete log-likelihood, with respect to the current conditional distribution of L given Z and the current estimates of the parameters $w^{(t)}$ and $\theta^{(t)}$:

$$Q(w,\theta|w^{(t)},\theta^{(t)}) = \mathbb{E}_{L|Z,w^{(t)},\theta^{(t)}} [\log L(w,\theta|z,l)]$$

We denote $u_{ic} = (F_{c_1}(z_i^1), ..., F_{c_Q}(z_i^Q))$ and $\xi_{ic} = (\Phi^{-1}(u_{ic}^1), ..., \Phi^{-1}(u_{ic}^Q)), 1 \leq i \leq n, c \in \mathcal{C}$. One has:

$$\begin{aligned} Q(w,\theta|w^{(t)},\theta^{(t)}) &= \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} \left[\log(w_c) + \sum_{q:c_q=0} \log(\phi(z_i^q)) + \sum_{q:c_q=1} \log(f_1^q(z_i^q)) + \log(c_\theta(u_{ic})) \right] \\ &= \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} \log(c_\theta(u_{ic})) + \text{cste} \\ &= \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} \log \left[|\theta|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \xi_{ic}^T (\theta^{-1} - I) \xi_{ic} \right) \right] + \text{cste} \\ &= \frac{n}{2} \log(|\theta^{-1}|) - \frac{1}{2} \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} (\xi_{ic}^T (\theta^{-1} - I) \xi_{ic}) + \text{cste} \\ &= \frac{n}{2} \log(|\theta^{-1}|) - \frac{1}{2} \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} (\xi_{ic}^T \theta^{-1} \xi_{ic}) + \text{cste} \\ &= \frac{n}{2} \log(|\theta^{-1}|) - \frac{1}{2} \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} \text{tr}\left(\xi_{ic}^T \theta^{-1} \xi_{ic} \right) + \text{cste} \\ &= \frac{n}{2} \log(|\theta^{-1}|) - \frac{1}{2} \sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} \text{tr}\left(\xi_{ic} \xi_{ic}^T \theta^{-1} \right) + \text{cste} \\ &= \frac{n}{2} \log(|\theta^{-1}|) - \frac{1}{2} \text{tr}\left(\sum_{i=1}^n \sum_{c\in\mathcal{C}} \hat{\tau}_{ic} (\xi_{ic} \xi_{ic}^T) \theta^{-1} \right) + \text{cste} \\ &= \frac{n}{2} \log(|\Lambda|) - \frac{1}{2} \text{tr}(S\Lambda) + \text{cste} \end{aligned}$$

where $\hat{\tau}_{ic} = \frac{w_c^{(t)} \psi^c_{\theta^{(t)}}(Z_i)}{\sum_{c \in \mathcal{C}} w_c^{(t)} \psi^c_{\theta^{(t)}}(Z_i)}$, $\Lambda = \theta^{-1}$ and $S = \sum_{i=1}^n \sum_{c \in \mathcal{C}} \hat{\tau}_{ic}(\xi_{ic}\xi_{ic}^T)$.
Differentiating the log-likelihood with respect to $\Lambda = \theta^{-1}$ leads to

$$\begin{aligned}\frac{\delta}{\delta\Lambda} \log L(\theta|x) &= \frac{n}{2}\frac{\delta}{\delta\Lambda}\log(|\Lambda|) - \frac{1}{2}\frac{\delta}{\delta\Lambda}\mathrm{tr}\,(S\Lambda) \\ &= \frac{n}{2}(\Lambda^{-1})^T - \frac{1}{2}S^T \\ &= \frac{n}{2}\theta^T - \frac{1}{2}\left(\sum_{i=1}^n \sum_{c \in C} \hat{\tau}_{ic}(\xi_{ic}\xi_{ic}^T)\right)^T\end{aligned}$$

Setting the derivative at 0, one obtains

$$\frac{\delta}{\delta\Lambda} \log L(\theta|x) = 0 \Leftrightarrow \hat{\theta} = \frac{1}{n}\sum_{i=1}^n \sum_{c \in C} \hat{\tau}_{ic}(\xi_{ic}\xi_{ic}^T) \tag{4}$$

Note that the explicit expression of $\widehat{\theta}$ obtained in Eq. (4) guarantees that the estimated matrix is symmetric positive-definite, but not necessarily with value 1 on the diagonal. In practice, one needs to check that $\hat{\theta}$ is properly a correlation matrix at each (M) step of the EM procedure, and whenever needed correct the initial estimate by scaling the resulting covariance matrix into a correlation matrix.

Supplementary Note 2: Equivalence between Benjamini-Hochberg correction on p-values and the estimation of lFDR from the posteriors

In the classical setting of the test procedure, one can control the FDR at a nominal level of α by applying the Benjamini-Hochberg (BH) procedure. Assume n test were performed and let $p_{(1)} \leq ... \leq p_{(n)}$ be the set of ordered p-values. Define the value r such as

$$r = \arg\max_k \left\{ k \,/\, \frac{n_0 p_{(k)}}{k} \leq \alpha \right\}.$$

The set of rejected hypotheses is then $S_{BH} = \{i_{(1)}, ..., i_{(r)}\}$, where $i_{(k)}$ is the index associated to p-value $p_{(k)}$. Here n_0 is the number of items for which $\mathcal{H}_0$ holds, which is usually unknown. In practice, n_0 can be upper bounded by n or adaptively obtained from the data, which will be designated hereafter by the adaptive BH procedure(16).

Alternatively, in the setting of a mixture model with two components 0 and 1 with emission distributions ψ_0, ψ_1 and priors W_0, W_1, one can define the posterior

$$\tau(z) = \frac{W_1 \psi_1(z)}{W_0 \psi_0(z) + W_1 \psi_1(z)},$$

The local FDR can then be approximately controlled at a nominal level of α by applying the following empirical procedure (12, 13). Assume n items were generated under the previous mixture model and let $\tau_{(1)} \leq ... \leq \tau_{(n)}$ be the set of ordered posteriors. Define the value m such as

$$m = \arg\max_k \left\{ k \,/\, \frac{1}{k} \sum_{1=1}^{k} (1 - \tau_{(i)}) \leq \alpha \right\}.$$

The set of items classified into class 1 is then $S_{lFDR} = \{i_{(1)}, ..., i_{(m)}\}$, where $i_{(k)}$ is the index associated to posterior $\tau_{(k)}$.

The connection between the two approaches in the context of CHT is established by the next proposition:

Proposition 1: Consider n observations drawn according to mixture model (1) of Section 2.2, and let $\hat{\tau}_i$ and $\widehat{\text{pval}}_i$ be the posterior and p-value associated to item i, respectively. Consider the adaptive BH procedure where n_0 is estimated by $n\hat{W}_0$, where $\hat{W}_0 = \sum_{c \in \mathcal{C}_0} \hat{w}_c$. Then $S_{BH} = S_{lFDR}$.

Proof:

For a given item i, the p-value corresponding to the composite hypothesis test is constructed as follows:

$$\widehat{\text{pval}}_i = \frac{1}{nW_0} \sum_{j=1}^{n} \mathbb{1}_{\{\hat{\tau}_j > \hat{\tau}_i\}} (1 - \hat{\tau}_j) \tag{5}$$

Let $\widehat{\text{pval}}_{(1)} < ... < \widehat{\text{pval}}_{(n)}$ represent the ordered set of p-values, and $i_{(k)}$ is the index associated to p-value $\widehat{\text{pval}}_{(k)}$.

According to the adaptive BH procedure, the estimate of the False Discovery Rate when the set of rejected hypotheses is $S_{BH} = \{i_{(1)}, ..., i_{(k)}\}$ is given by:

$$\widehat{FDR}(k) = \frac{\hat{n}_0}{k} \widehat{\text{pval}}_{(k)} \tag{6}$$

with $\hat{n}_0$ an estimate of the number of true null hypotheses. Here, one can estimate the number of null hypotheses as $\hat{n}_0 = nW_0$. Hence, we have:

$$\begin{aligned} \widehat{FDR}(k) &= \frac{nW_0}{k} \widehat{\text{pval}}_{(k)} \\ &= \frac{nW_0}{k} \frac{1}{nW_0} \sum_{j=1}^{n} \mathbb{1}_{\{\hat{\tau}_j > \hat{\tau}_k\}} (1 - \hat{\tau}_j) \\ &= \frac{1}{k} \sum_{j=1}^{n} \mathbb{1}_{\{\hat{\tau}_j > \hat{\tau}_k\}} (1 - \hat{\tau}_j) \end{aligned}$$

By the definition of the p-value 5, $\widehat{\text{pval}}_i < \widehat{\text{pval}}_j \Leftrightarrow \hat{\tau}_j > \hat{\tau}_i$.

Consequently:

$$\begin{aligned}\widehat{FDR}(k) &= \frac{1}{k}\sum_{j=1}^{n}\mathbb{1}_{\{\widehat{\text{pval}}_j<\widehat{\text{pval}}_{(k)}\}}(1-\hat{\tau}_j)\\ &= 1-\frac{1}{k}\sum_{j=1}^{n}\mathbb{1}_{\{\hat{\tau}_j>\hat{\tau}_k\}}\hat{\tau}_j\\ &= \widehat{lFDR}(\hat{\tau}_k)\end{aligned}$$

where $\widehat{lFDR}(\hat{\tau}_k)$ is the estimate of the false discovery rate when thresholding at the level $\hat{\tau}_k$ defined by (12) and (13). Thus, applying the adaptive BH correction with $\hat{n}_0 = nW_0$ to p-value Eq. (5) is equivalent to using the local FDR rejection rule defined by (12) and (13). ■

Supplementary Note 3: Algorithm for the Memory-Efficient EM

Here we present how the burden associated to the storage of the posterior matrix can be avoided in the EM algorithm. Quantities in red correspond to the ones that are actually stored in place of the posterior matrix for later use.

Algorithm 1 Proportions w_c and copula parameter θ inference using EM with memory handling

Inputs: f0Mat $\in \mathcal{M}_{n\times Q}$ containing the $f_0(Z_i^q)$.
f1Mat $\in \mathcal{M}_{n\times Q}$ containing the $\hat{f}_1^q(Z_i^q)$.
$\pi_0 \in \mathbb{R}^Q$ containing the proportions $\hat{\pi}_0^q$.

$\hat{w}_c^{(t)} \leftarrow \prod_{q:c_q=0} \pi_0^q \prod_{q:c_q=1}(1-\pi_0^q)$ ▷ initialisation
$\hat{\theta}^{(t)} \leftarrow$ Kendall's tau estimate

while $||\hat{w}_c^{(t)} - \hat{w}_c^{(t+1)}|| > \epsilon$ or $||\hat{\theta}^{(t)} - \hat{\theta}^{(t+1)}|| > \epsilon$ **do**
  Memory-Efficient step E:
  for i in $1:n$ **do**
    $S_i^{(t)} \leftarrow 0$
    for $c \in \mathcal{C}$ **do**
      $S_i^{(t)} \leftarrow S_i^{(t)} + \hat{w}_c^{(t)} \psi^c_{\hat{\theta}^{(t)}}(Z_i)$ ▷ computation of $\psi^c_{\hat{\theta}^{(t)}}(Z_i)$ on the fly
    end for
  end for

  Step M:
  for $c \in \mathcal{C}$ **do**
    $\hat{w}_c^{(t+1)} \leftarrow 0$
    for i in $1:n$ **do**
      $\hat{w}_c^{(t+1)} \leftarrow \hat{w}_c^{(t+1)} + \frac{\hat{w}_c^{(t)} \psi^c_{\hat{\theta}^{(t)}}(Z_i)}{S_i^{(t)}}$ ▷ computation of $\psi^c_{\hat{\theta}^{(t)}}(Z_i)$ on the fly
    end for
    $\hat{w}_c^{(t+1)} \leftarrow \frac{1}{n}\hat{w}_c^{(t+1)}$
  end for

  $\hat{\theta}^{(t+1)} \leftarrow 0$
  for i in $1:n$ **do**
    for $c \in \mathcal{C}$ **do**
      $\hat{\theta}^{(t+1)} \leftarrow \hat{\theta}^{(t+1)} + \hat{\tau}_{ic}(\xi_{ic}\xi_{ic}^T)$ ▷ computation of $\psi^c_{\hat{\theta}^{(t)}}(Z_i)$ on the fly
    end for
  end for
  $\frac{1}{n}\hat{\theta}^{(t+1)} \leftarrow \hat{\theta}^{(t+1)}$
  $t \leftarrow t+1$
end while

Outputs: $\hat{w}_c \in \mathbb{R}^{2^Q}$ containing the configuration proportions estimates.
$\hat{\theta}$ the estimate of the copula parameter.

Supplementary Note 4: Details for generating each simulation setting

Table 6. Description of the 12 settings considered for the simulation study.

Setting	Q	Correlation	Scenario
1	2	0.3	sparse
2			dense
3		0.5	sparse
4			dense
5	8	0.3	sparse
6			dense
7		0.5	sparse
8			dense
9	16	0.3	sparse
10			dense
11		0.5	sparse
12			dense

Each dataset consists of a $n \times Q$ matrix of Z-scores generated as follows. Markers are clustered into $K+1$ clusters. Each cluster is characterized by a configuration c, i.e. all markers of cluster k share the same configuration $c_k = (c_{k1}, ..., c_{kQ})$. One has

$$Z_i \sim \mathcal{N}(2c_k, \Sigma_\rho),$$

where Σ_ρ is a $Q \times Q$ correlation matrix where all covariances between traits are set to ρ. Assuming that matrix Σ_ρ is common to all clusters, a given cluster k is characterized by its associated number of markers and associated configuration vector. We considered different values for Q (2,8,16), two scenarios: $sparse$ or $dense$, and two levels of correlation ($\rho = 0.3, 0.5$). Denoting $n_{50,\tilde{q}}$ and $n_{100,\tilde{q}}$ the number of clusters of 50 and 100 markers respectively, which are associated with exactly $\tilde{q}$ randomly chosen traits, the details for generating a dataset for each simulation setting are presented below.

Settings 1 to 4: $Q = 2$***.***

- Sparse scenario: $n_{50,1} = n_{50,2} = 20$ and $n_{100,1} = n_{100,2} = 10$, such that $w_{00} = 0.96$, $w_{01} + w_{10} = 0.02$ and $w_{11} = 0.02$.

- Dense scenario: $n_{50,1} = 400$, $n_{100,1} = 200$ and $n_{50,2} = 100$, $n_{100,2} = 50$, such that $w_{00} = 0.5$, $w_{01} + w_{10} = 0.4$ and $w_{11} = 0.1$.

Settings 5 to 8: $Q = 8$***.***

- Sparse scenario: $n_{50,j} = 20$ and $n_{100,j} = 10$ for $j \in \{1,2,4,6,8\}$, such that $w_0 = 0.90$ for the markers being associated with none of the traits and $w_j = 0.02$ for the markers being associated with exactly $j \in \{1,2,4,6,8\}$ traits.

- Dense scenario: $n_{50,j} = 100$ and $n_{100,j} = 50$ for $j \in \{1,2,4,6,8\}$, such that $w_0 = 0.50$ for the markers being associated with none of the traits and $w_j = 0.1$ for the markers being associated with exactly $j \in \{1,2,4,6,8\}$ traits.

Settings 9 to 12: $Q = 16$***.***

- Sparse scenario: $n_{50,q} = 20$ and $n_{100,q} = 10$ for $q \in \{1,2,4\}$, $n_{50,\tilde{q}} = 10$ and $n_{100,\tilde{q}} = 5$ for $\tilde{q} \in \{6,8,14,16\}$, such that $w_0 = 0.90$ for the markers being associated with none of the traits, $w_q = 0.02$ for the markers being associated with exactly $q \in \{1,2,4\}$ traits, and $w_{\tilde{q}} = 0.01$ for the markers being associated with exactly $\tilde{q} \in \{6,8,14,16\}$ traits.

- Dense scenario: $n_{50,q} = 100$ and $n_{100,q} = 50$ for $q \in \{1,2,4\}$, $n_{50,\tilde{q}} = 50$ and $n_{100,\tilde{q}} = 25$ for $\tilde{q} \in \{6,8,14,16\}$, such that $w_0 = 0.5$ for the markers being associated with none of the traits, $w_q = 0.1$ for the markers being associated with exactly $q \in \{1,2,4\}$ traits, and $w_{\tilde{q}} = 0.05$ for the markers being associated with exactly $\tilde{q} \in \{6,8,14,16\}$ traits.

Supplementary Note 5: Additional simulations results

Table 7. False discovery rate and detection power of the procedures in sparse and dense scenarios with $Q = 2$ and $\rho = 0.5$ (settings 3 and 4).

	Sparse		Dense	
Method	FDR	Power	FDR	Power
DACT_Efron	0 (0)	0 (0)	0 (0)	0 (0)
DACT_JC	0 (0)	0 (0.001)	0.026 (0.018)	0.01 (0.006)
DACT_NO	**0.104 (0.033)**	0.092 (0.022)	0.055 (0.013)	0.042 (0.008)
HDMT	**0.41 (0.017)**	0.34 (0.011)	**0.374 (0.015)**	0.53 (0.021)
PLACO	**0.3 (0.083)**	0.007 (0.002)	0.052 (0.025)	0.011 (0.003)
adaFilter	**0.217 (0.02)**	0.185 (0.014)	**0.121 (0.007)**	0.146 (0.006)
qch	**0.437 (0.015)**	0.387 (0.013)	**0.239 (0.006)**	0.299 (0.006)
IMIX	**0.871 (0.021)**	0.848 (0.063)	**0.295 (0.079)**	0.448 (0.081)
c-csmGmm	0 (0)	0 (0)	0 (0)	0 (0)
Primo	**0.081 (0.022)**	0.057 (0.005)	**0.112 (0.009)**	0.139 (0.005)
qch_copula	0.057 (0.025)	0.04 (0.007)	0.061 (0.008)	0.067 (0.005)

For each procedure the FDR and the detection power (averaged over 25 runs) are displayed. Numbers in brackets correspond to standard errors. False discovery rate above 0.05 are in boldface.

Table 8. Detection power of the procedures for sparse and dense scenarios with $Q = 2$, $\rho = 0.3$ (settings 1 and 2).

Method	Sparse	Dense
DACT_Efron	0 (0.001)	0 (0)
DACT_JC	0.002 (0.002)	0.083 (0.021)
DACT_NO	0.027 (0.013)	0.006 (0.004)
HDMT	0.264 (0.017)	0.166 (0.006)
PLACO	0.011 (0.005)	0.012 (0.004)
adaFilter	0.123 (0.011)	0.094 (0.005)
qch	0.307 (0.018)	0.18 (0.005)
IMIX	0.471 (0.237)	0.248 (0.027)
c-csmGmm	0 (0)	0 (0)
Primo	0.057 (0.006)	0.085 (0.004)
qch_copula	0.054 (0.009)	0.051 (0.005)

For each procedure the detection power (averaged over 25 runs) are displayed. Numbers in brackets correspond to standard errors.

Table 9. Performance of the procedures Primo, adaFilter and qch_copula for the simulation settings 7-8 ($Q = 8$) and 11-12 ($Q = 16$) with $\rho = 0.5$.

			Primo		adaFilter		qch_copula	
Q	Scenario	Test	FDR	Power	FDR	Power	FDR	Power
8	Dense	at least 2	0.005 (0.001)	0.144 (0.004)	0.03 (0.002)	0.315 (0.003)	0.058 (0.002)	0.617 (0.004)
		at least 4	0.006 (0.001)	0.147 (0.005)	0.02 (0.002)	0.199 (0.003)	0.064 (0.003)	0.549 (0.005)
		at least 8	0.075 (0.006)	0.178 (0.006)	0.043 (0.005)	0.111 (0.008)	0.104 (0.007)	0.28 (0.006)
	Sparse	at least 2	0.008 (0.003)	0.115 (0.01)	0.126 (0.008)	0.29 (0.008)	0.063 (0.005)	0.422 (0.007)
		at least 4	0.01 (0.004)	0.12 (0.011)	0.076 (0.008)	0.196 (0.006)	0.067 (0.007)	0.415 (0.011)
		at least 8	0.142 (0.031)	0.09 (0.01)	0.078 (0.019)	0.107 (0.009)	0.104 (0.023)	0.119 (0.012)
16	Dense	at least 2	-	-	0.034 (0.002)	0.265 (0.003)	0.057 (0.002)	0.722 (0.004)
		at least 4	-	-	0.025 (0.002)	0.182 (0.003)	0.067 (0.002)	0.756 (0.004)
		at least 8	-	-	0.028 (0.003)	0.139 (0.004)	0.083 (0.002)	0.681 (0.005)
	Sparse	at least 2	-	-	0.153 (0.008)	0.246 (0.006)	0.08 (0.007)	0.647 (0.007)
		at least 4	-	-	0.105 (0.008)	0.18 (0.006)	0.094 (0.008)	0.724 (0.009)
		at least 8	-	-	0.072 (0.008)	0.138 (0.008)	0.1 (0.008)	0.582 (0.014)

For each procedure the FDR and the detection power (averaged over 25 runs) are displayed. Numbers in brackets correspond to standard errors.

Supplementary Note 6: Information on the datasets used for the application cases studies

A. Application I: Detection of pleiotropic regions associated to psychiatric disorders. We analyzed 14 psychiatric disorders obtained from the Psychiatric Genomics Consortium (PGC) [1]. The disorders studied are anorexia nervosa (AN), anxiety disorder (AD), attention-deficit/hyperactivity disorder (ADHD), autism spectrum disorder (ASD), alcohol use disorders (AUDIT-T alcohol use traits identification test based on total score, AUDIT-C alcohol use traits identification test based on consumption, AUDIT-P alcohol use traits identification test based on problematic consequences of drinking), bipolar disorder (BIP), cannabis use (CU[2]), major depression disorder (MDD), obsessive-compulsive disorder (OCD), post-traumatic stress disorder (PTSD), schizophrenia (SCZ) and Tourette's syndrome (TS). Table 10 details the information about each study. For all the disorders, we obtained their European-only summary statistics and performed the following quality control:
(i) excluded non-biallelic SNPs and those with strand-ambiguous alleles;
(ii) removed duplicated SNPs.

Table 10. Summary information of 14 psychiatric disorders GWAS analyzed in this study.

Traits	N(case/control)	PGC study	Reference
AN	14,477 (3,495/10,982)	an2017	(38)
AD	17,526 (5,761/11,765)	anx2016	(39)
ADHD	53,293 (19,099/34,194)	adhd2019	(40)
ASD	46,350 (18,381/27,969)	asd2019	(41)
AUDIT-T	121,604	sud2019-alcuse	(42)
AUDIT-C	121,604	sud2019-alcuse	(42)
AUDIT-P	121,604	sud2019-alcuse	(42)
BIP	51,710 (20,352/31,358)	bip2019	(43)
CU	184,765 (53,180/131,585)	-	(44)
MDD	480,359 (135,458/344,901)	mdd2018	(45)
OCD	9,725 (2,688/7,037)	ocd2018	(46)
PTSD	174,659 (23,212/151,447)	ptsd2019	(47)
SCZ	77,096 (33,640/43,456)	scz2019asi	(48)
TS	14,307 (4,819/9,488)	ts2019	(49)

B. Application II: Detection of viruses resistance hotspots regions in cucumber. Our second application case is based on GWAS summary statistics derived from the experiment described in (27). In this study, a panel of 226 elite cucumber lines from Bayer germplasm collections, 40 landraces from Bayer's internal collections, and 23 hybrids-including resistance controls-were inoculated with six viruses to evaluate their responses. The viruses included zucchini yellow mosaic virus (ZYMV), papaya ringspot virus (PRSV), watermelon mosaic virus (WMV), cucumber mosaic virus (CMV), cucumber vein yellowing virus (CVYV), and cucumber green mottle mosaic virus (CGMMV). An individual GWAS was conducted on a number of SNPs ranging from $n = 378,049$ to $n = 424,393$ depending on the virus. The experimental protocol and results are fully detailed in (27). The aim of the study was to identify QTLs associated with virus resistance in cucumber.

[1] as available from the PGC website: https://pgc.unc.edu/for-researchers/download-results/
[2] available from the following website: https://drive.google.com/drive/folders/1uFBeUwE-0UA0YmLTVQ9Ym_oU2HkHPdhk

Supplementary Note 7: Supplementary figures for applications

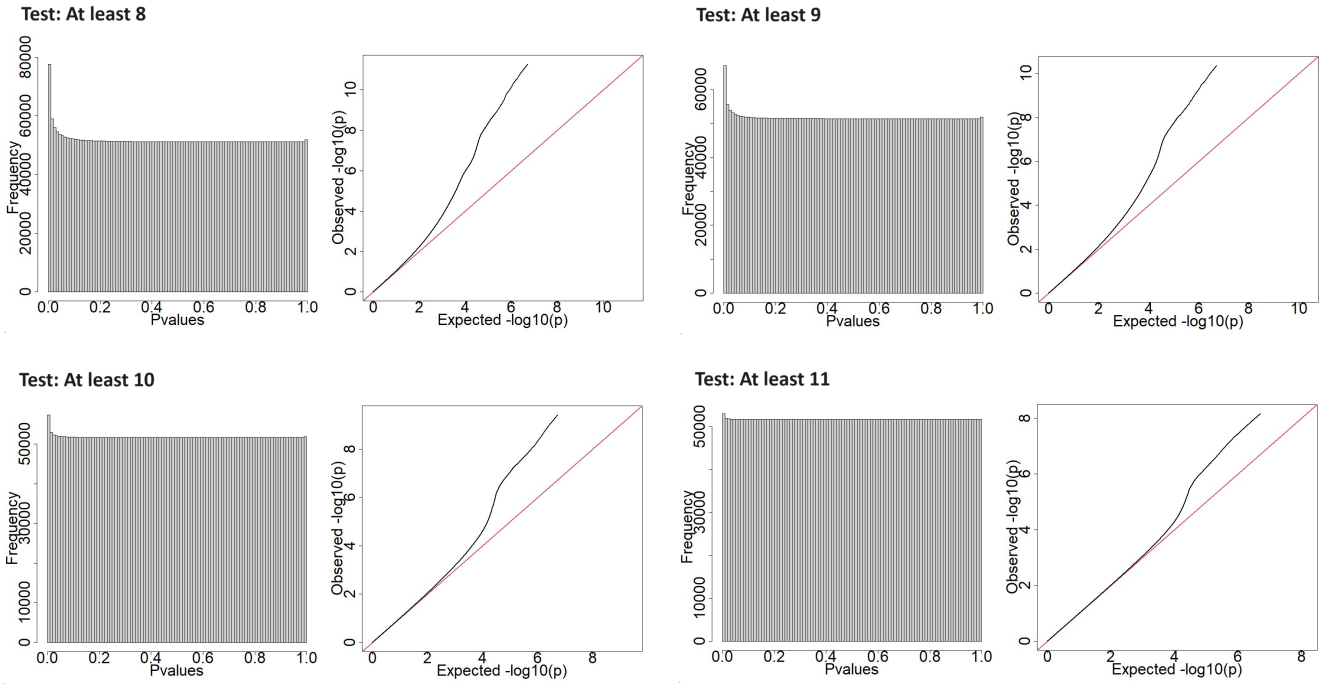

Fig. 3. Quality control of the p-values distribution corresponding to the different composite hypothesis tests applied on the 14 psychiatric disorders GWAS.
Left: Histogram of the qch_copula p-values. **Right:** Quantile-Quantile plot of the qch_copula p-values.

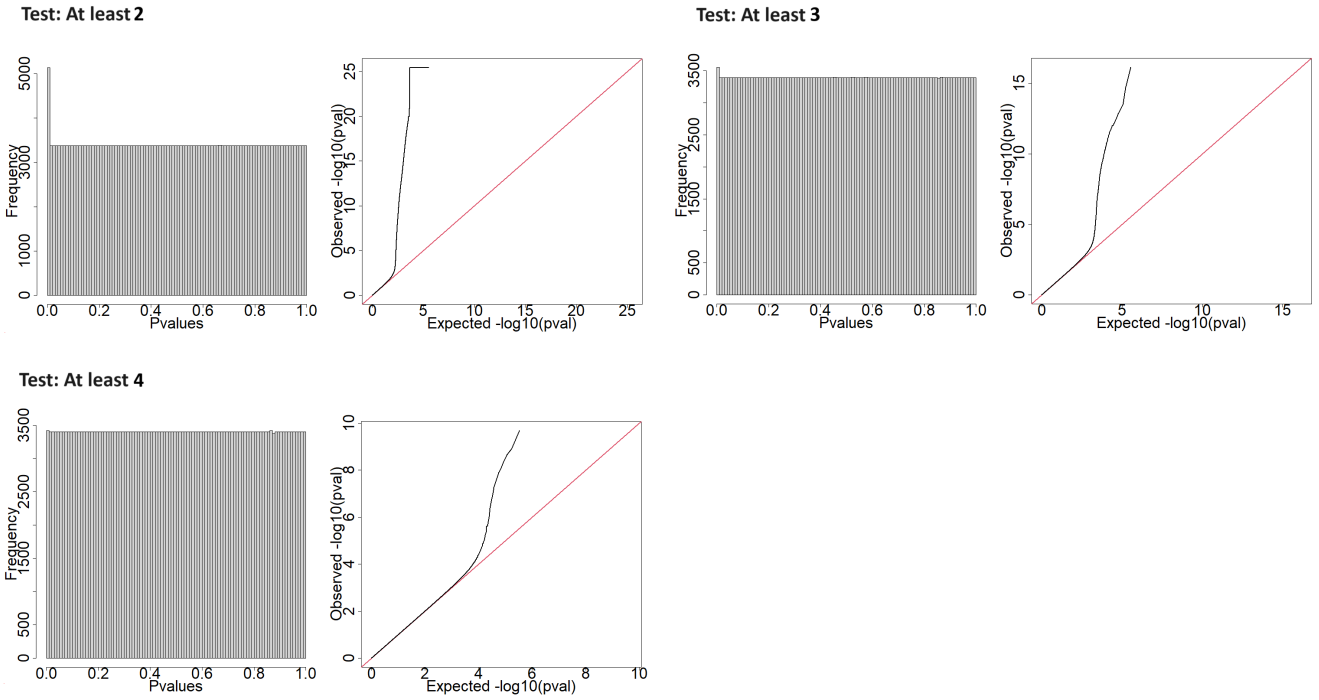

Fig. 5. Quality control of the p-values distribution corresponding to the different composite hypothesis tests applied on the six viruses resistance GWAS.
Left: Histogram of the qch_copula p-values. **Right:** Quantile-Quantile plot of the qch_copula p-values.

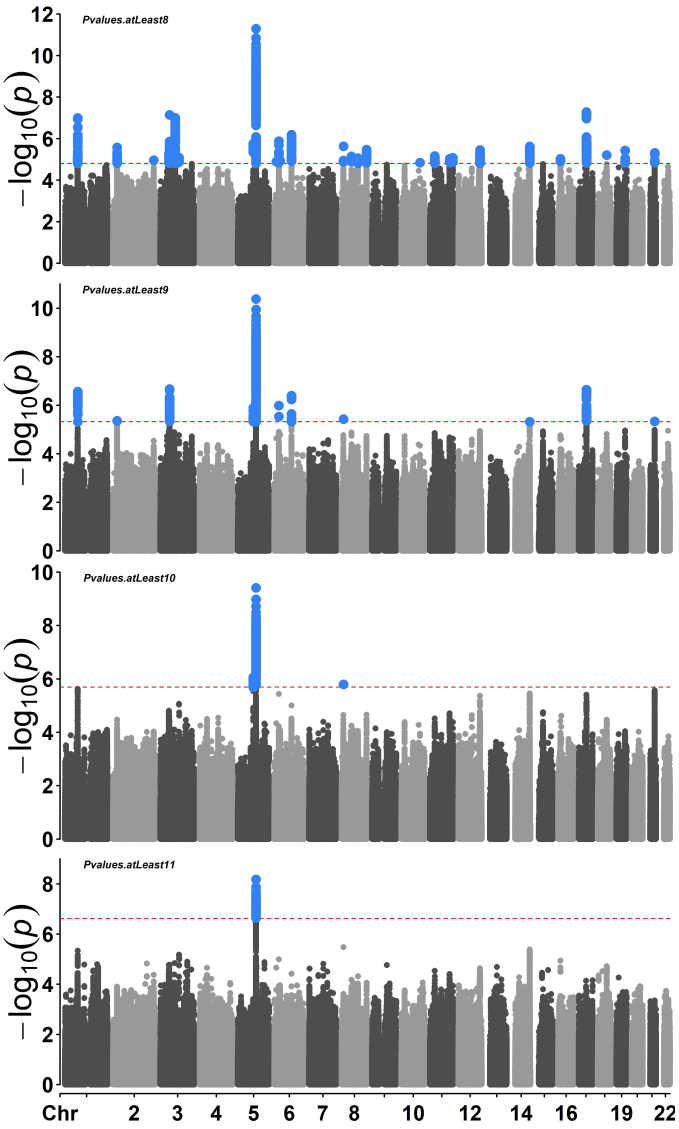

Fig. 4. -log10(p-values) of the different composite hypothesis tests along the chromosomes applied on the 14 psychiatric disorders GWAS.

From top to bottom, the composite hypothesis test identifies the SNPs associated with at least 8, 9, 10 and 11 disorders, respectively. The red dotted line represents the significance threshold at a nominal false discovery rate of 0.05. Significant SNPs are represented in blue.

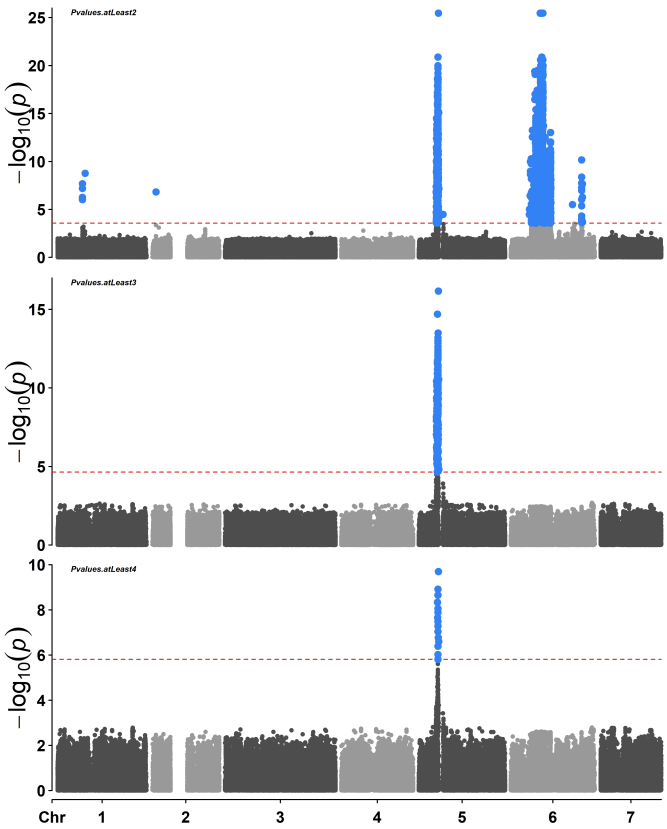

Fig. 6. -log10(p-values) of the different composite hypothesis tests along the chromosomes applied on the six viruses resistance GWAS.

From top to bottom, the composite hypothesis test identifies the SNPs associated with the resistance to at least 2, 3 and 4 viruses, respectively. The red dotted line represents the significance threshold at a nominal false discovery rate of 0.05. Significant SNPs are represented in blue.