



ODBAE: a high-performance model identifying complex phenotypes in high-dimensional biological datasets



Yafei Shen^{1,2,7}, Tao Zhang^{3,7}, Zhiwei Liu⁴, Kalliopi Kostelidou⁵, Ying Xu^{④,6}✉ & Ling Yang^{①,2}✉

Identifying complex phenotypes from high-dimensional biological data is challenging due to the intricate interdependencies among different physiological indicators. Traditional approaches often focus on detecting outliers in single variables, overlooking the broader network of interactions that contribute to phenotype emergence. Here, we introduce ODBAE (Outlier Detection using Balanced Autoencoders), a machine learning method designed to uncover both subtle and extreme outliers by capturing latent relationships among multiple physiological parameters. ODBAE's revised loss function enhances its ability to detect two key types of outliers: influential points (IP), which disrupt latent correlations between dimensions, and high leverage points (HLP), which deviate from the norm but go undetected by traditional autoencoder-based methods. Using data from the International Mouse Phenotyping Consortium (IMPC), we show that ODBAE can identify knockout mice with complex, multi-indicator phenotypes—normal in individual traits, but abnormal when considered together. In addition, this method reveals novel metabolism-related genes and uncovers coordinated abnormalities across metabolic indicators. Our results highlight the utility of ODBAE in detecting joint abnormalities and advancing our understanding of homeostatic perturbations in biological systems.

Phenotypes, including symptoms, defined as the observable characteristics or traits of an organism, are the result of complex interactions between genotype and environmental factors^{1,2}. These traits are often quantified using physiological or pathological indicators that serve as proxies for the underlying biological or pathological processes. Recent advances in phenotypic research, particularly through genetic manipulation, have deepened our understanding of gene function and the mechanisms underlying pathological conditions across a wide range of organisms, including plants and animals³. Traditionally, gene-related phenotypic analyses have focused on identifying abnormalities in individual physiological indicators^{3–5}. However, findings from the International Mouse Phenotyping Consortium (IMPC) demonstrate strong correlations between physiological indicators, suggesting that many traits and diseases result not from isolated abnormalities, but from coordinated disruptions across multiple physiological indicators⁶.

Our current understanding of phenotypes or diseases often emphasizes the detection of outliers in individual physiological indicators, typically

those outside the normal physiological range. However, the emergence of disease in an organism is a more complex process. Homeostasis—the dynamic balance maintained by biological systems—can be perturbed at multiple levels before a single-indicator deviates outside the normal range⁷. This suggests that phenotypic abnormalities may manifest as imbalances between correlated indicators, even when each individual measure remains within its expected range. These imbalances may represent early warning signs of disease or dysfunction that may be missed by traditional univariate analysis.

We hypothesize that the coordinated perturbation of physiological indicators reflects an emerging phenotype or disease state, even when individual measures appear normal. By examining these subtle interdependencies, we can gain insight into how homeostasis is perturbed in knockout mouse models. Furthermore, for genes whose knockout does not result in obvious outliers in individual parameters, disruption of homeostatic balance may still occur through correlated indicators. This approach challenges the notion that the absence of a detectable abnormal phenotype

¹School of Mathematical Sciences, Soochow University, Suzhou, Jiangsu, China. ²Center for Systems Biology, Soochow University, Suzhou, Jiangsu, China.

³Jiangsu Province Engineering Research Center of Development and Translation of Key Technologies for Chronic Disease Prevention and Control, Suzhou Vocational Health College, Suzhou, Jiangsu, China. ⁴Cambridge-Suda Genomic Research Center, Suzhou Medical College of Soochow University, Suzhou, Jiangsu, China. ⁵KTL Research Acceleration LTD, Nicosia, Cyprus. ⁶Jiangsu Key Laboratory of Neuropsychiatric Diseases, Suzhou Medical College of Soochow University, Suzhou, Jiangsu, China. ⁷These authors contributed equally: Yafei Shen, Tao Zhang. ✉e-mail: yingxu@suda.edu.cn; lyang@suda.edu.cn

in knockout mice indicates no functional impact, and instead suggests that our focus on single-indicator abnormalities may obscure more complex systemic effects.

Existing methods, such as linear regression^{6,8} and basic bioinformatics tools^{9–11}, have provided initial insights into these relationships. However, these techniques struggle to capture non-linear and complex relationships between multiple physiological indicators, especially in high-dimensional biological datasets. Machine learning methods, particularly autoencoders, have emerged as powerful tools for detecting outliers by learning complex, latent patterns within data^{12–18}. By compressing and reconstructing data, autoencoders can capture subtle variations and relationships that may indicate phenotypic abnormalities across multiple indicators. Their ability to learn from high-dimensional data makes them well suited for analyzing the intricate dynamics involved in phenotypic expression^{19,20}. In the regression analysis, all outliers can be categorized as influential points (IP) and high leverage points (HLP). IP exerts a substantial influence on the model's fitting results and HLP deviate from the center of the dataset but may not necessarily impact the model fit²¹. Since the reconstruction of any dataset by the autoencoder can be regarded as a process of regression, IP can be easily detected when using autoencoders for outlier detection. However, there are still many uncertainties about the detection effect of HLP.

In this study, we introduce Outlier Detection using Balanced Auto-encoders (ODBAE), a machine learning method designed to detect outliers by capturing the relationships between physiological indicators. ODBAE improves on traditional autoencoder by balancing the reconstruction error across principal component directions, enhancing its ability to detect both IP that deviate from the expected relationships between indicators and HLP that are far from the data center. We apply ODBAE to developmental and metabolic datasets from IMPC to explore how coordinated perturbations in physiological parameters can reveal previously unidentified phenotypes, even in cases where knockout mice show no abnormalities in individual traits. By using ODBAE, we aim to identify novel gene functions and phenotypic complexities that have been missed by single-indicator screening methods. Collectively, these findings establish ODBAE as a powerful tool for identifying complex anomalies and unknown phenotypes within biological systems.

Results

Outlier detection in ODBAE

ODBAE takes tabular datasets from various sources as input, such as gene knockout mouse dataset from IMPC and patient health examination data. Each row in the data represents a record, and each column represents an attribute. The goal of ODBAE is to maintain the effectiveness of the auto-encoder in detecting IP while further improving the detection performance for HLP. This is a comprehensive approach to outlier detection in high-dimensional tabular datasets. Furthermore, ODBAE also provides an anomaly explanation for each outlier, indicating which parameter or parameters are associated with the anomaly. For gene knockout mouse datasets, ODBAE can identify genes that, when knocked out, lead to anomalies, as well as the abnormal parameter pairs associated with the gene knockout.

Overview of ODBAE

To address the challenge of outlier detection in high-dimensional tabular datasets, we construct an improved autoencoder model (ODBAE) with refined training loss function. Our model leverages the strong representation learning and feature extraction capabilities of traditional autoencoders for IP, while mitigating their limitations in HLP detection ("Methods"). The core principle underlying ODBAE-based outlier detection is that inliers can be well reconstructed, whereas outliers, including both HLP and IP, may generate significant reconstruction errors. ODBAE performs unsupervised outlier detection on the input high-dimensional datasets. During training process, ODBAE learns as much intrinsic information as possible from training set and reconstructs the training dataset by minimizing the loss function. The training dataset selection strategy depends on the initial

outlier prevalence within the dataset. For datasets with few outliers presence, the entire dataset are utilized for both training and testing for outlier detection, such as data from wild-type mice in IMPC. However, if the overall proportion of outliers are not clear, a subset of the data exhibiting fewer anomalies is chosen for model training. Then, the entire dataset, or other subset, can be used for outlier detection. Subsequently, the trained model is applied to the test dataset, generating reconstruction error of all sample points. Finally, sample points with reconstruction errors greater than the predefined threshold are considered outliers. As shown in Fig. 1a, outlier detection from biological datasets by ODBAE include three steps: (1) given a training dataset as input, the intrinsic information of normal data points are learned; (2) the trained model is used to reconstruct the test dataset, and outliers will be identified according to large reconstruction error ("Methods"); (3) the detected outliers including HLP and IP are explained using highest reconstruction errors and kernel-SHAP to gain the abnormal parameters²².

Based on mathematical analysis, ODBAE uses a revised training loss function incorporating an appropriate penalty term to Mean Square Error (MSE) to balance the reconstruction by properly suppressing complete reconstruction of the autoencoder ("Methods"). Specifically, the penalty term can ensure the equal eigenvalue difference between each principal component direction of the training and reconstructed dataset. Therefore, the incorporated penalty term enhance the detection of HLP, while the MSE term maintain the identification performance of IP.

To explain each outlier detected by ODBAE, we first identify the top features that contribute the most to the reconstruction error, and then apply kernel-SHAP to obtain the features that have the greatest impact on them (Fig. 1b and "Methods"). Finally, ODBAE outputs instance-based outliers and their explanations. For outliers in categories, if the anomaly rate for a category of outliers is greater than the set threshold, then ODBAE provide anomaly explanation for each category according to their mean values of each feature.

ODBAE's comprehensive benefits

The power of ODBAE lies in its ability to efficiently detect both HLP and IP in complex, high-dimensional biological data sets. HLP are outliers that deviate from the central distribution, while IP significantly influence model results. To evaluate the performance of ODBAE in detecting IP, we compared it to principal component analysis²³, a common outlier detection method. ODBAE demonstrated superior performance, particularly in identifying IP, due to its enhanced ability to capture nonlinear relationships between data points (Supplementary Fig. 1a–d).

We then applied ODBAE to phenotypic data from the IMPC to identify complex phenotypes in knockout mouse models. Using a threshold of the top 2% of absolute z-scores²⁴ for any given physiological parameter, we flagged mice as potential outliers ("Methods"). We then applied ODBAE with a predefined abnormality ratio (Fig. 1c). If more than 50% of mice from a single-gene knockout strain were classified as outliers, the corresponding gene was identified as significant and further analyzed (Fig. 1d). ODBAE was able to identify key mutant strains with complex phenotypes, even in cases where individual physiological indicators appeared normal.

To illustrate, we analyzed eight developmental parameters: Body length (BL), Body weight (BW), Bone Area (BA), Bone Mineral Density (excluding skull), Distance travelled total, Forelimb and hindlimb grip strength measurement mean, heart rate, and Heart weight. Using wild-type mice as the training set and knockout mice as the test set, we evaluated 1904 single-gene knockout mouse strains ("Methods"). Given the prevalence of sexual dimorphism in disease phenotypes^{3,25}, males and females were analyzed separately (Supplementary Data 1a–d). In the female dataset, ODBAE identified *Ckb* null mice as outliers despite their individual parameter values being within the normal range. These mice had normal BL and BW, but their body mass index (BMI) was abnormally low. Specifically, four of the eight *Ckb* null mice had extremely low BMI values, calculated as $BMI = BW (g)/BL^2 (cm)$ ^{26–28}. Analysis revealed that while BL and BW were within the expected ranges, their relationship was abnormal, leading to the

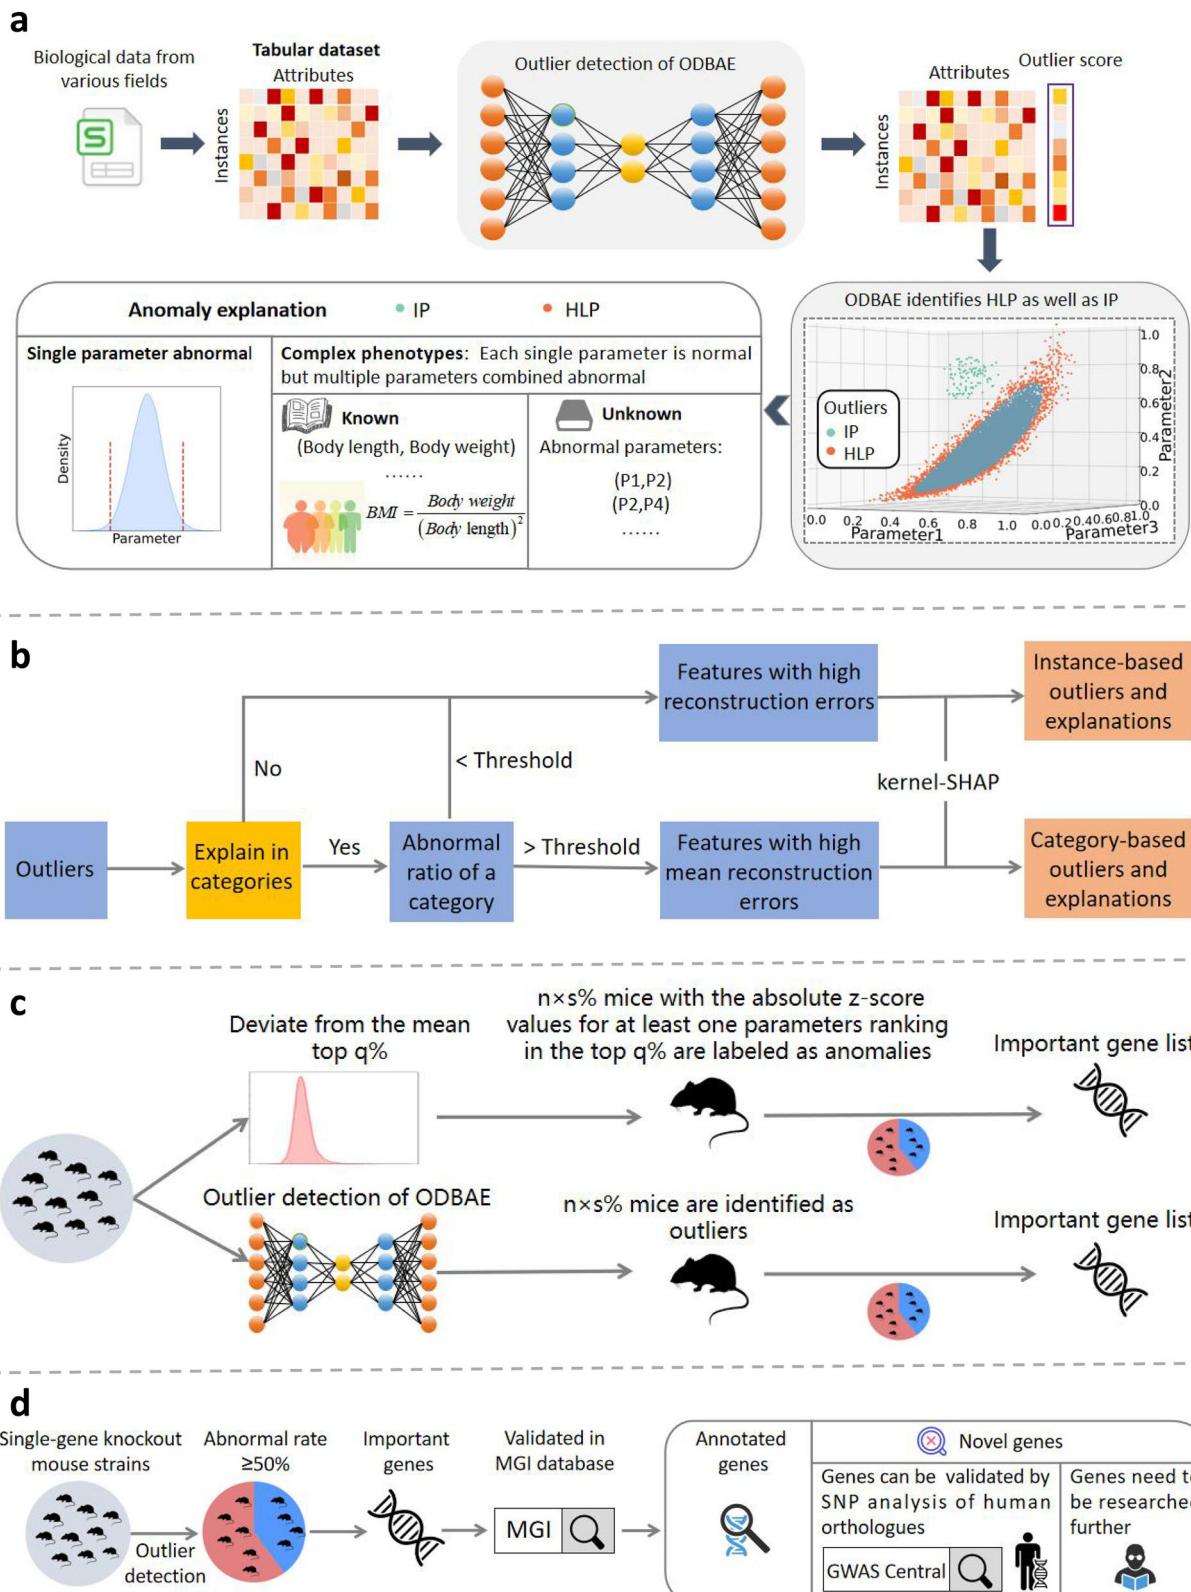
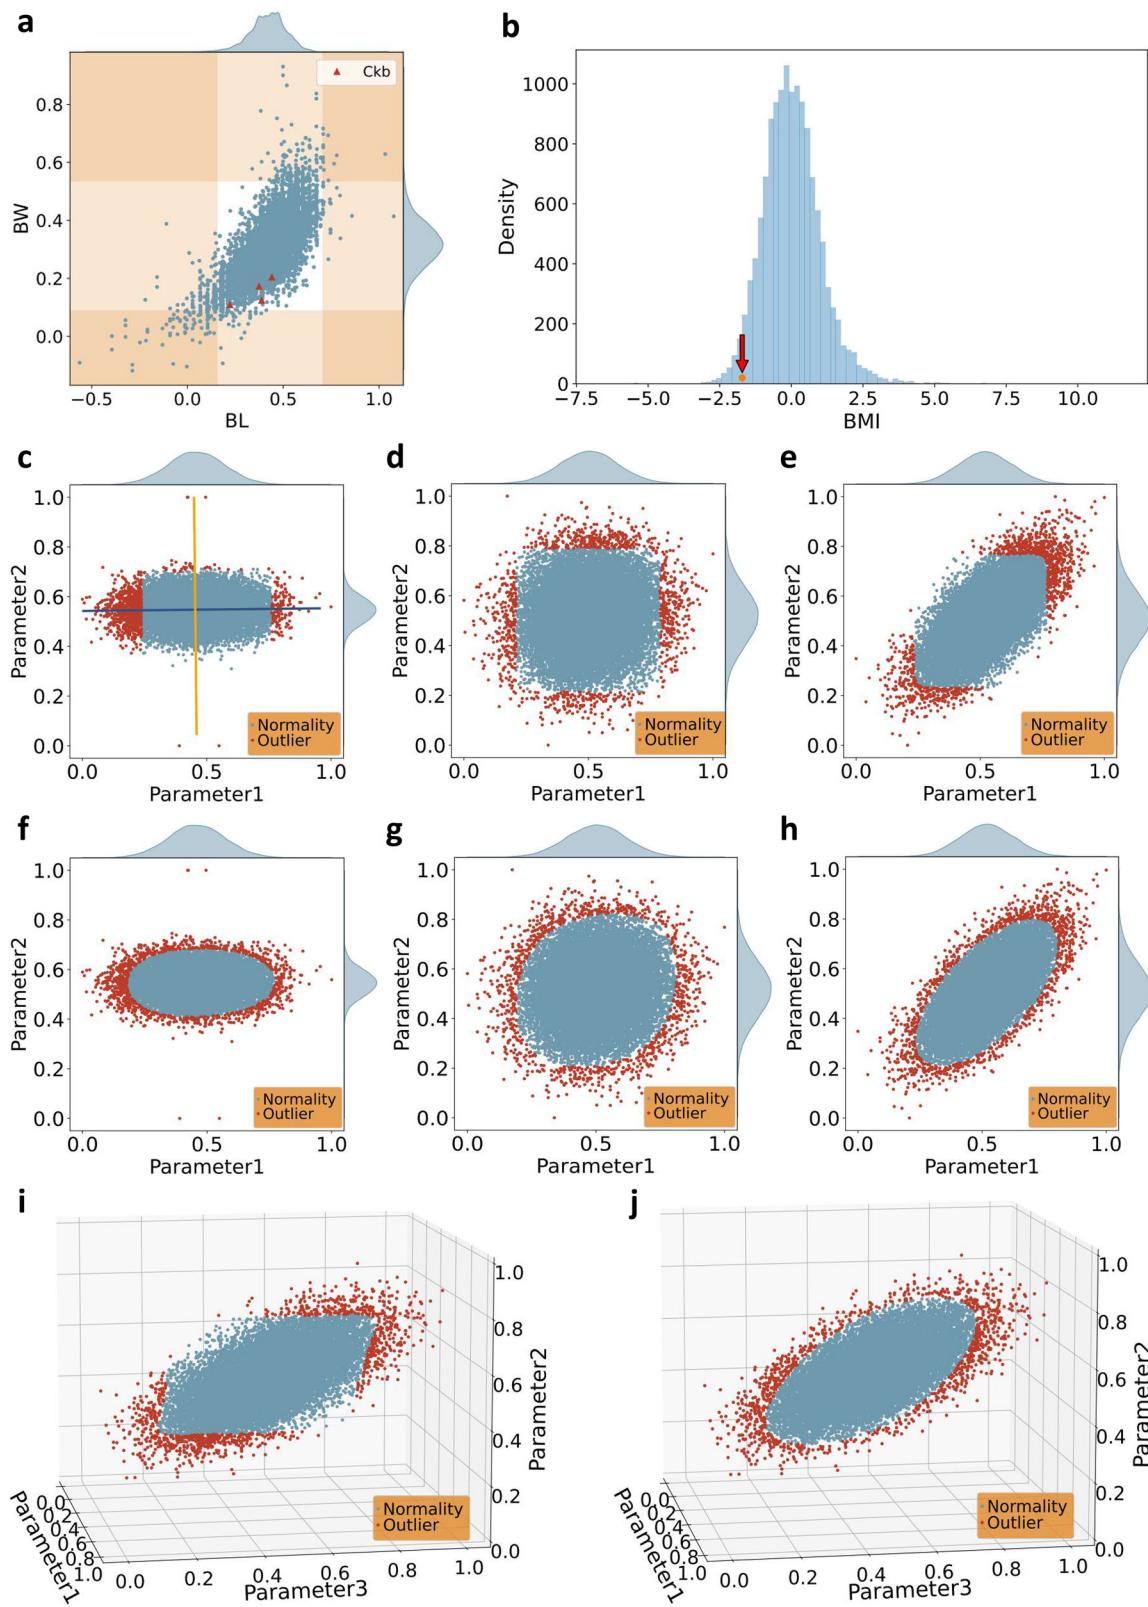


Fig. 1 | ODBAE is an approach for identifying and explaining outliers in tabular datasets, it can discover complex phenotypes and identify novel genes when applied to knockout mouse datasets. a Process for discovering complex phenotypes. ODBAE first detect outliers based on a balanced autoencoder, which can balance the reconstruction of the training dataset to detect both HLP and IP, then perform anomaly explanation to gain abnormal parameters. Finally, some outliers are identified due to joint anomalies of multiple parameters. **b** Anomaly explanation of ODBAE. The detected outliers are explained in terms of SHAP values. **c** For both

males and females, abnormal mice are first identified according to the presence of metabolic parameters that deviated too far from the mean, and then the same abnormal proportion is set to identify abnormal mice with ODBAE. Finally, some important genes will only be obtained by ODBAE. **d** Genes with an abnormal proportion of more than 50% of the corresponding single-gene knockout mouse strains are considered important genes by ODBAE. If an important gene has no known metabolic links in the MGI database, it will be considered a novel gene and further validated by SNP analysis of human orthologues.



identification of these mice as outliers (Fig. 2a). The average BMI of these *Ckb*-deficient mice was lower than 97.14% of other mice (Fig. 2b), demonstrating the sensitivity of ODBAE to complex, multivariate outliers. In addition, *Ckb* has previously been implicated in developmental processes and obesity, further strengthening the biological significance of these findings²⁹.

ODBAE also identified previously unknown phenotypes, highlighting its potential to uncover novel gene functions. In terms of the identification of HLP, most studies choose MSE as the training loss function and force the autoencoder to completely reconstruct the training dataset, which brings two major limitations: (1) only HLP along principal component directions with poor reconstruction can be detected; (2) in the process of anomaly

Fig. 2 | ODBAE accurately identify outliers with multi-dimensional joint anomalies. **a** Scatter plot of parameters BL and BW for mouse developmental dataset, the data points corresponding to mice with the gene *Ckb* knockout deviate significantly from most of the data points, but their values of both BL and BW are within the normal range. The BL or BW values of the data points in the shaded part are ranked in the top 2% based on the absolute value of the z-score. **b** The distribution of standardized BMI of all knockout mice in mouse developmental dataset, the mean BMI of *Ckb* knockout mice is significantly smaller. **c–h** Outlier detection results for 3 synthetic 2-dimensional Gaussian distribution datasets when the outlier ratio is

0.05, including detection results of MSE-trained autoencoder and ODBAE for 2-dimensional Gaussian distribution dataset with diagonal covariance matrix and non-Gaussian distributed noise (**c, f**), 2-dimensional Gaussian distribution dataset with diagonal covariance matrix (**d, g**) and 2-dimensional Gaussian distribution dataset with non-diagonal covariance matrix (**e, h**). The orange line and grey line in (**c**) represent two principal directions. **i, j** Outlier detection results of MSE-trained autoencoder (**i**) and ODBAE (**j**) for 3-dimensional dataset with an intrinsic dimension of 2 when the outlier ratio is 0.05, and the data points follow a 2-dimensional Gaussian distribution with respect to Parameter1 and Parameter3.

detection, each dimension of the dataset is independently detected. To evaluate the performance difference between our modified loss function and MSE in HLP detection, we used these two trained autoencoder to detect outliers for several datasets following multi-dimensional Gaussian distributions. For example, data in Fig. 2c–e were three 2-dimensional Gaussian distribution datasets and data in Fig. 2i was 3-dimensional dataset whose intrinsic dimension was two, the data points followed a 2-dimensional Gaussian distribution with respect to Parameter1 and Parameter3 (“Methods”). The detection result in Fig. 2c shows that most of the detected HLP lie in the direction of the principal components indicated by the blue line, but few lie in the other direction. In Fig. 2d, HLP in two dimensions of the dataset are independently detected, making outliers that are anomalous in multiple dimensions (outliers at the four corners) undetectable. Fig. 2i demonstrates the same issue in three-dimensional space. Moreover, these issues persisted when the activation function was replaced with another function (Supplementary Figs. 2a–d and 3a–d). This is because the saturating nature of the non-linear activation function causes the autoencoder to have different reconstruction capabilities for data points located at the boundaries of the dataset compared to those located in the interior (Supplementary Fig. 4a, c and Supplementary Note 1). Usually, the distribution of a dataset along various principal component directions may not be balanced, and the number of data points falling into unreconstructable regions may also be uneven (“Methods”). This leads to different detection performances for HLP along different principal component directions. Besides, the distribution of reconstructed dataset is independent across all dimensions (Supplementary Fig. 4a, c). Therefore, the new loss function in ODBAE is designed to ensure that the reconstructed dataset follows a balanced joint distribution (Supplementary Fig. 4b, d). Then, we used ODBAE to detect outliers in these datasets and found that most outliers were identified (Fig. 2f–h, j).

We further evaluated ODBAE’s detection of both IP and HLP using a synthetic 3-dimensional dataset (Supplementary Fig. 5a). Area Under Curve (AUC) and Average Precision (AP) scores were used to evaluate the accuracy of outlier detection results (“Methods”). The generated 3-dimensional training dataset formed a 2-dimensional manifold, where the subspace represented by Parameter1 and Parameter3 followed a 2-dimensional Gaussian distribution (“Methods”, Supplementary Fig. 5a). Here, we assumed that the ratio of HLP and the ratio of IP were equal. During the anomaly detection, we used the training dataset as the test dataset. Then, we considered the top $[\delta n]$ data points in the subspace represented by Parameter1 and Parameter3, ranked by Mahalanobis distance, as HLP, where δ was the outlier ratio, and n was the number of the data points in the training dataset. Besides, we additionally generated $[\delta n]$ points that were not on the manifold of the training dataset and treated them as IP (Supplementary Fig. 5b). The results showed that ODBAE worked better than other anomaly detection methods (Supplementary Fig. 5c, d).

In summary, ODBAE is a powerful tool for detecting complex phenotypes in high-dimensional biological datasets, capturing both isolated and coordinated abnormalities. This capability allows the identification of novel gene functions and previously undetected phenotypes, particularly in cases where traditional methods fail to detect subtle perturbations in physiological homeostasis.

ODBAE outperforms other methods by a large margin

To evaluate the strengths of ODBAE, we performed outlier detection of ODBAE on low-dimensional synthetic datasets, high-dimensional synthetic datasets, and two benchmark datasets. ODBAE was compared with the autoencoder with MSE loss function (MSE-AE), autoencoder with Mean Absolute Error loss function (MAE-AE), and Deep Autoencoding Gaussian Mixture Model (DAGMM)³⁰. The comparison results were presented through AUC and AP scores.

We generated three datasets with 2-dimensional Gaussian distribution, and their intrinsic dimension were also 2. In one of the datasets, the correlation between the two dimensions was small (Fig. 3a); the correlation between the two dimensions in another dataset was relatively large (Fig. 3b); in the last dataset, the correlation between the two dimensions was also small and non-Gaussian distributed noise existed (Fig. 3c). Therefore, the outliers in these three datasets were all HLP. Then we set different outlier ratios and calculated the corresponding AUC and AP scores of the outlier detection results. To be specific, if the outlier ratio is δ and the number of sample points is n , then the input data is sorted according to their Mahalanobis distance, and the first $[\delta n]$ sample points are regarded as positive. For each fixed anomaly ratio δ , outlier detection was repeated 10 times on the dataset, and the average accuracy scores for each method were recorded. Fig. 3d–f show the AUC scores of test results on these three datasets, and Fig. 3g–i are corresponding AP scores, respectively. ODBAE outperformed other schemes in detecting HLP in all three datasets at any outlier ratio (Supplementary Data 2a–c).

Since autoencoders are often used for outlier detection in high-dimensional datasets, we also analyzed the outlier detection results of ODBAE on two high-dimensional synthetic datasets. Both datasets followed multi-dimensional Gaussian distribution and their covariance matrices were both diagonal matrices. So the intrinsic dimension of these two datasets were equal to their actual dimension. One of the datasets was 50-dimensional and the other dataset was 100-dimensional. Finally, ODBAE still performed best in HLP detection on high-dimensional datasets (Supplementary Fig. 6a–d).

To further evaluate the robustness of ODBAE under more realistic conditions where data distributions deviate from Gaussian assumptions, we applied the method to two additional synthetic datasets: one sampled from a two-dimensional Laplace distribution and the other from a two-dimensional uniform distribution. In both scenarios, ODBAE again exhibited superior outlier detection performance (Supplementary Fig. 7a–d).

Additionally, we tested ODBAE on two benchmark datasets: Dry Bean Dataset and Breast Cancer Dataset. Dry Bean Dataset contained 13611 instances and 17 attributes, and there were 7 classes in the original dataset. In each experiment, the inliers were 2000 data points from a certain class and $\delta \times 2000$ data points were randomly selected from the remaining 6 classes of data points as outliers, where δ was the outlier ratio and $\delta \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. For this dataset, we set the dimension of the hidden layer of the autoencoder to 8 during the experiment. Breast Cancer Dataset contained 569 instances and 31 attributes. All data points fell into two categories: benign and malignant. We randomly selected 300 data points from the data

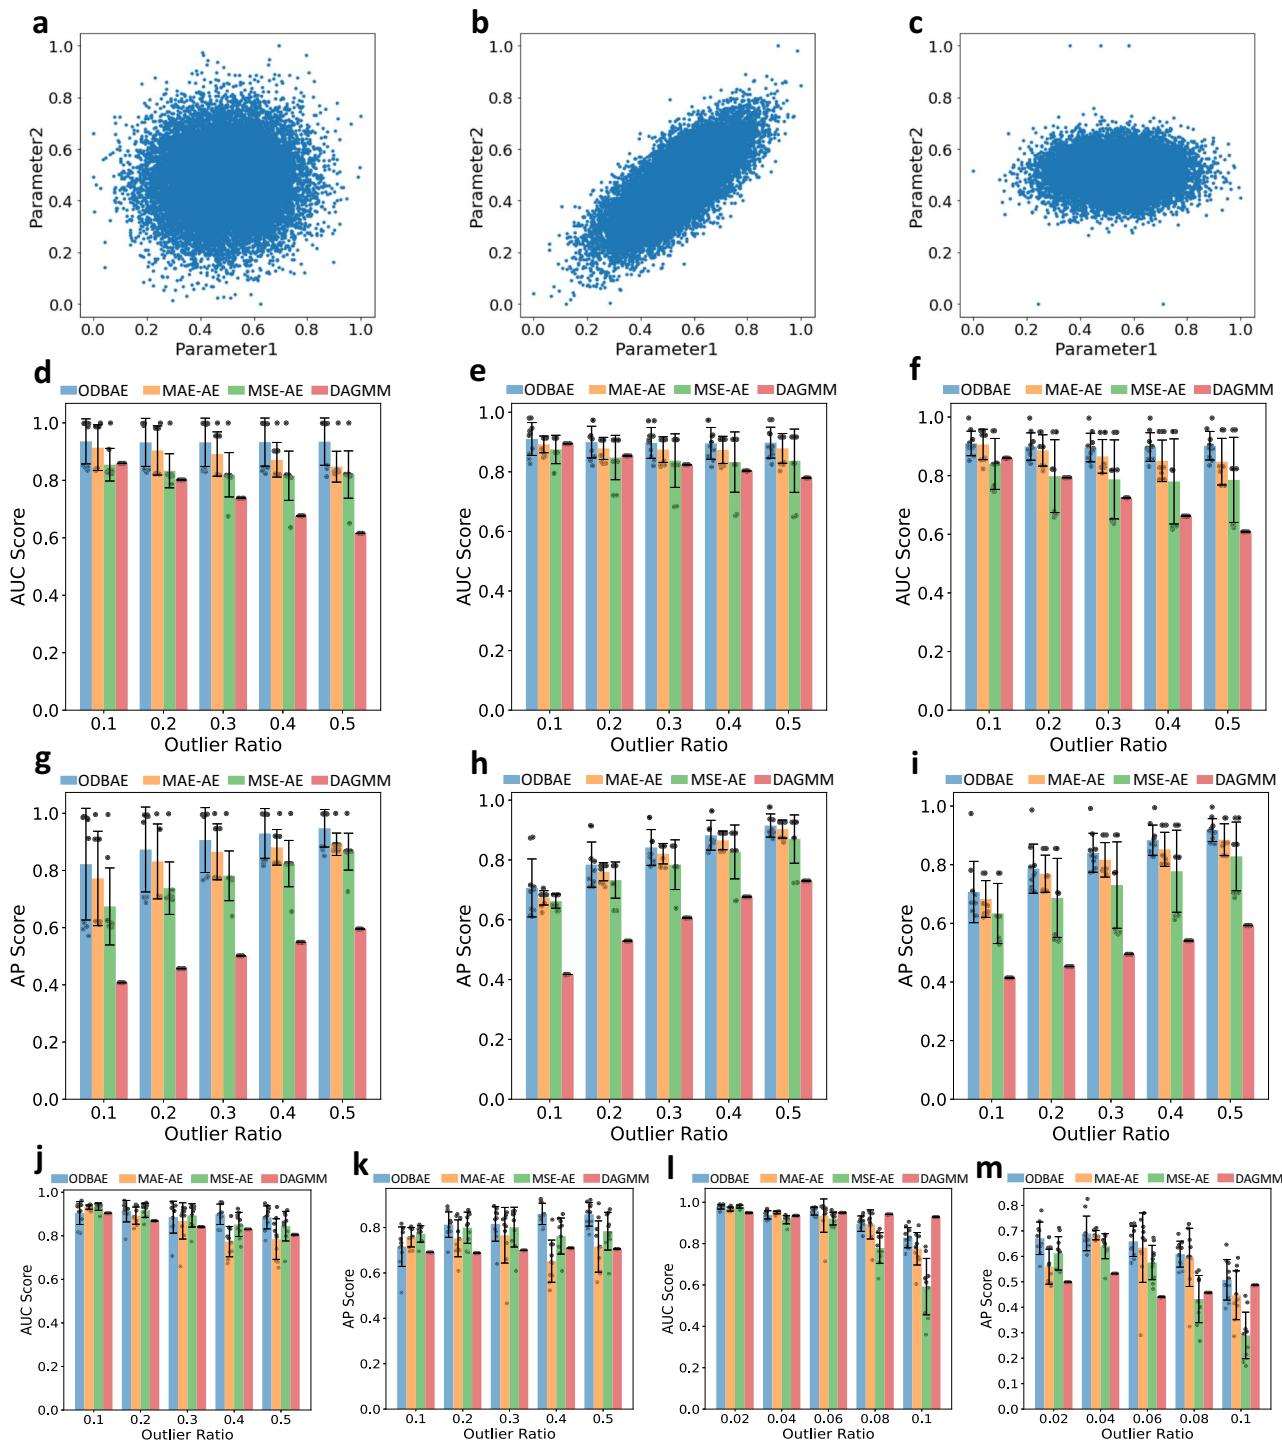


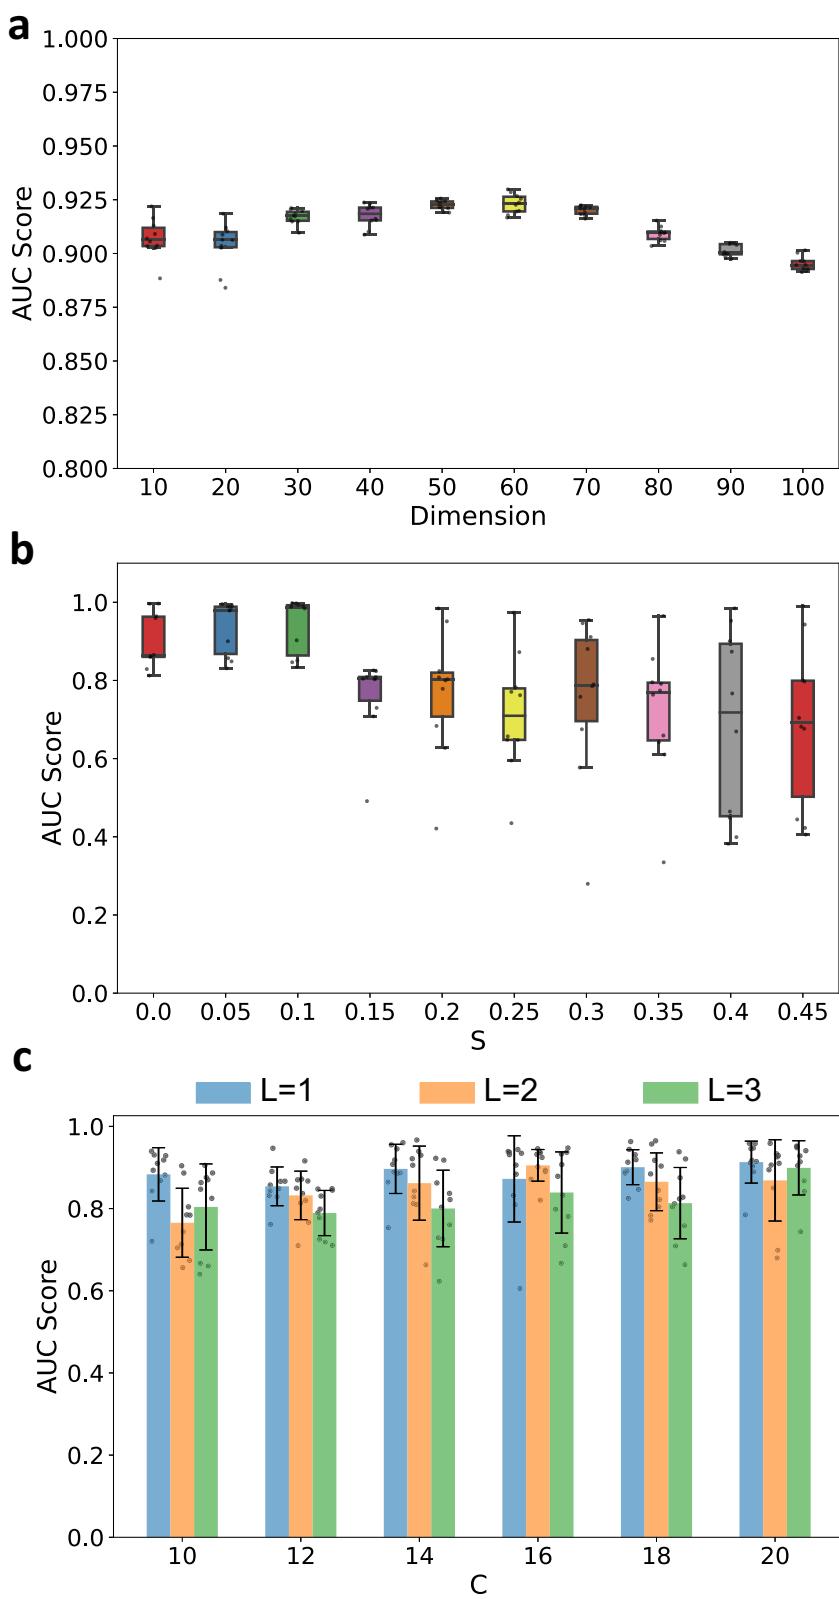
Fig. 3 | Outlier detection effect of ODBAE is generally better than other schemes. a–c Data point distribution plot for 3 synthetic dataset, including 2-dimensional Gaussian distribution dataset with diagonal covariance matrix (a), 2-dimensional Gaussian distribution dataset with non-diagonal covariance matrix (b) and 2-dimensional Gaussian distribution dataset with diagonal covariance matrix and non-Gaussian distributed noise (c). d–f AUC scores of comparison results of dataset

in this figure (a, d), (b, e), and (c, f). g–i AP scores of comparison results of dataset in this figure (a, g), (b, h), and (c, i). j, k Comparison results for Dry Bean Dataset. l, m Comparison results for Breast Cancer Dataset. In (d–m), all methods were run 10 times under different outlier ratios, and error bars indicate mean \pm standard error.

points labeled benign as inliers and $\delta \times 300$ data points from the other categories as outliers. For this dataset, the outlier ratio $\delta \in \{0.02, 0.04, 0.06, 0.08, 0.1\}$ and the dimension of the hidden layer of the autoencoder was 15. Overall, due to the improvement in HLP detection, ODBAE worked best on both benchmark datasets (Fig. 3j–m and Supplementary Data 2d, e).

In summary, ODBAE outperforms several existing autoencoder-based methods in detecting anomalies from both human-induced and benchmark datasets, indicating that ODBAE can more comprehensively identify anomalous situations. Therefore, ODBAE provides strong support for finding unknown anomalies and exploring unknown phenotypes in existing biological data.

Fig. 4 | ODBAE exhibits a certain degree of robustness. **a** Outlier detection performance of ODBAE across datasets with varying dimensionalities (Supplementary Data 2f). For each fixed-dimensional dataset, the model was trained 10 times, and the AUC score for outlier detection was recorded after each training run. **b** Impact of varying noise intensities on the performance of ODBAE (Supplementary Data 2g). Laplace-distributed noise $L(0, s)$ was added to the dataset to simulate different noise levels. For each fixed value of s , the model was trained 10 times on the noise-augmented data, and the AUC score for outlier detection was recorded after each training run. **c** Outlier detection accuracy of ODBAE with architecture (L, C) on the Dry Bean Dataset (Supplementary Data 2h). For each fixed architecture (L, C) , the model was trained 10 times, and the average AUC score across the 10 runs was recorded to evaluate performance. Error bars indicate mean \pm standard error.



ODBAE exhibits a certain degree of robustness

The performance of ODBAE, like any machine learning framework, is influenced by several hyperparameters and data conditions. We evaluated the impact of dataset dimensionality, noise levels, and model architecture on the outlier detection performance of ODBAE.

We generated Gaussian-distributed datasets of varying dimensionality to train ODBAE, and evaluated its outlier detection performance by measuring the AUC scores when the outlier ratio was set to 0.1. As illustrated in Fig. 4a, the anomaly detection performance of ODBAE remains largely stable as the dimensionality of the dataset increases. This observation

highlights the potential applicability of ODBAE to outlier detection tasks in high-dimensional datasets.

Considering that real-world datasets frequently contain noise, we further evaluated the impact of data perturbations on the anomaly detection performance of ODBAE. A two-dimensional Gaussian-distributed dataset V containing 10,000 samples was generated. To simulate varying noise intensities, we added perturbation ϵ to each data point, where ϵ was drawn from a Laplace distribution $L(0, s)$, with larger s values representing higher noise levels. We then evaluated the performance of ODBAE trained on the noisy datasets in detecting outliers within the original dataset V , with the outlier ratio set to 0.1. As shown in Fig. 4b, ODBAE maintained high outlier detection accuracy under moderate noise levels ($s \leq 0.1$), highlighting its robustness in the presence of noise typically observed in real-world data.

In addition, we investigated the influence of model architecture on the performance of ODBAE. We similarly employed the Dry Bean dataset, randomly selecting 2000 samples from a single class to form the training dataset. Additionally, 200 outlier samples were randomly drawn from the remaining 6 classes. We then evaluated the outlier detection accuracy of ODBAE with varying model architectures denoted as (L, C) , where the model consists of $2L + 1$ hidden layers, with all layers except the bottleneck layer having dimensionality C . The results demonstrate that both the number of layers L and the dimensionality of each layer C significantly influence the performance of ODBAE (Fig. 4c). Therefore, in practical applications, tuning these architectural parameters—specifically, the number of layers and the dimensionality of each layer—can help optimize the outlier detection performance of ODBAE.

In summary, ODBAE maintains high outlier detection accuracy in high-dimensional datasets and exhibits robustness to moderate levels of noise. Furthermore, its performance can be further optimized in practical applications by tuning the number of layers and the dimensionality of each layer in the model architecture.

ODBAE identifies new metabolic genes

We applied ODBAE to metabolism-related datasets from IMPC to uncover genetic elements involved in metabolic regulation. Fourteen key metabolic parameters were selected for analysis, including: Alanine aminotransferase (ALA), Albumin (Alb), Alkaline phosphatase, Aspartate aminotransferase (ASA), Calcium (Ca), Creatinine (Cre), Glucose (Glu), HDL-cholesterol (HDLC), Phosphorus (Ph), Total bilirubin (TB), Total cholesterol (TC), Total protein (TP), Triglycerides (TG), and Urea. In total, we analyzed 45922 mice from 3064 single-gene knockout strains, including 23024 males and 22898 females.

ODBAE was used to detect outliers in both male and female data sets. Initially, mice with absolute z-scores in the top 1.2% for at least one metabolic parameter were flagged as outliers. Of the 23024 male mice, 2672 were flagged as abnormal, and of the 22898 female mice, 2553 abnormalities were identified. To ensure robust detection, we adjusted the proportion of anomalies and identified the top 10% of mice with the highest reconstruction errors as outliers. Genes were considered significant if more than 50% of mice from a given knockout strain were identified as outliers, suggesting a strong association between these genes and metabolism. In total, ODBAE identified 128 significant genes in males (9.5%) and 111 significant genes in females (8.08%), for a total of 203 genes (Fig. 5a, b and Supplementary Data 1e–h).

Moreover, we observed that the metabolism-associated genes identified by ODBAE exhibited pronounced sex-specific differences between male and female mice (Fig. 5b). To evaluate the biological significance of this sexual dimorphism, we focused on metabolism-related genes exhibiting pronounced sexual dimorphism. Specifically, we selected genes for which both male and female mouse strains included at least seven individuals, and for which the proportion of knockout strains displaying abnormal phenotypes exceeded 50% in only one sex (either male or female), while remaining at 0% in the other. In total, 18 genes were identified as showing significant sexual dimorphism, 10 of which had been previously reported as sex-

dimorphic in earlier studies (Fig. 5c, “Methods” and Supplementary Data 3i). These findings confirm the sexually dimorphic nature of these genes and underscore the utility of ODBAE as a robust tool for uncovering metabolism-associated genetic mechanisms.

Notably, 91 of these 203 genes could not be detected using traditional z-score based methods. This indicates that the metabolic abnormalities associated with these genes involve complex, multi-parameter phenotypes. ODBAE’s anomaly detection includes a detailed explanation of outliers based on the highest reconstruction error and kernel-SHAP analysis, which allows us to identify the specific metabolic parameters driving the anomalies. For example, knockout of the *Crabp1* gene resulted in abnormal levels of Alb and HDLC, even though the individual z-scores for these parameters were within normal ranges. Visualization of these data (Fig. 5d) shows that while the z-scores for Alb and HDLC appeared normal, the knockout mice deviated significantly from the central distribution of all data points, highlighting the complex metabolic perturbation caused by *Crabp1* knockout. More examples are shown in Supplementary Fig. 8a–d.

Further investigation using the Mouse Genome Informatics (MGI) database revealed that 43 of the 203 genes (21.18%) identified by ODBAE had no previously known association with metabolic phenotypes (Fig. 5e and Supplementary Data 3a). To explore the potential relevance of these newly identified genes, we analyzed their human orthologues (Supplementary Data 3b). Single-nucleotide polymorphisms (SNPs) within $a \pm 1$ kb region of these orthologues were extracted from GWAS Central³¹, resulting in a total of 804 SNPs (Supplementary Data 3c). We then investigated their association with metabolic diseases, focusing on type 2 diabetes (T2D). Specifically, for each SNP, we evaluated the extent of association across T2D-related traits on data from the DIAGRAM, GIANT, and GLGC consortia^{32–34}. Cross-phenotype meta-analysis (CPMA)³⁵ identified SNPs in four gene regions—*TWF2*, *TMED10*, *HOXA10*, and *NBAS*—that were strongly linked to T2D-related traits (CPMA $p < 0.05$) (Supplementary Data 3d). Visualization of the corresponding outliers confirmed that these genes exhibited distinct patterns of deviation across key metabolic dimensions (Supplementary Fig. 9a–d). In addition to T2D, many of these novel genes may be associated with other metabolic disorders. KEGG pathway analysis (via KOBAS) for mouse datasets revealed significant enrichment in metabolic pathways³⁶, including glycolysis, propanoate metabolism, and pyruvate metabolism (Fig. 5f and Supplementary Data 3e). These results suggest that ODBAE not only identifies genes involved in complex metabolic phenotypes but also pinpoints pathways that are disrupted by genetic perturbations.

Overall, ODBAE provides a powerful tool for uncovering novel metabolic genes and phenotypes that are undetectable using traditional z-score-based methods. It successfully identifies genes associated with complex, multi-parameter metabolic abnormalities and reveals novel genetic contributors to metabolic diseases, many of which were previously unannotated in the MGI database.

ODBAE integrates novel metabolic phenotypes

ODBAE’s ability to detect outliers based on joint abnormalities, beyond traditional z-score methods, highlights its strength in integrating complex metabolic phenotypes. Here, we demonstrate how ODBAE can detect metabolic abnormalities in gene knockout mice by identifying abnormal relationships between multiple metabolic parameters, rather than focusing on single-parameter outliers. Our analysis of metabolism-related datasets from IMPC reveals that a large proportion of outliers (89.14% of male and 92.75% of female mice) exhibit abnormalities in multiple metabolic parameters simultaneously. These findings suggest that joint abnormalities, or perturbations in the correlations between parameters, are important for detecting phenotypic changes in knockout mice. For example, knockout of the *Ckb* gene disrupts the relationship between BW and BL, resulting in an abnormal BMI, even though BW and BL are individually within normal ranges. This highlights how gene knockouts often disrupt intrinsic correlations between physiological parameters, resulting in complex metabolic phenotypes.

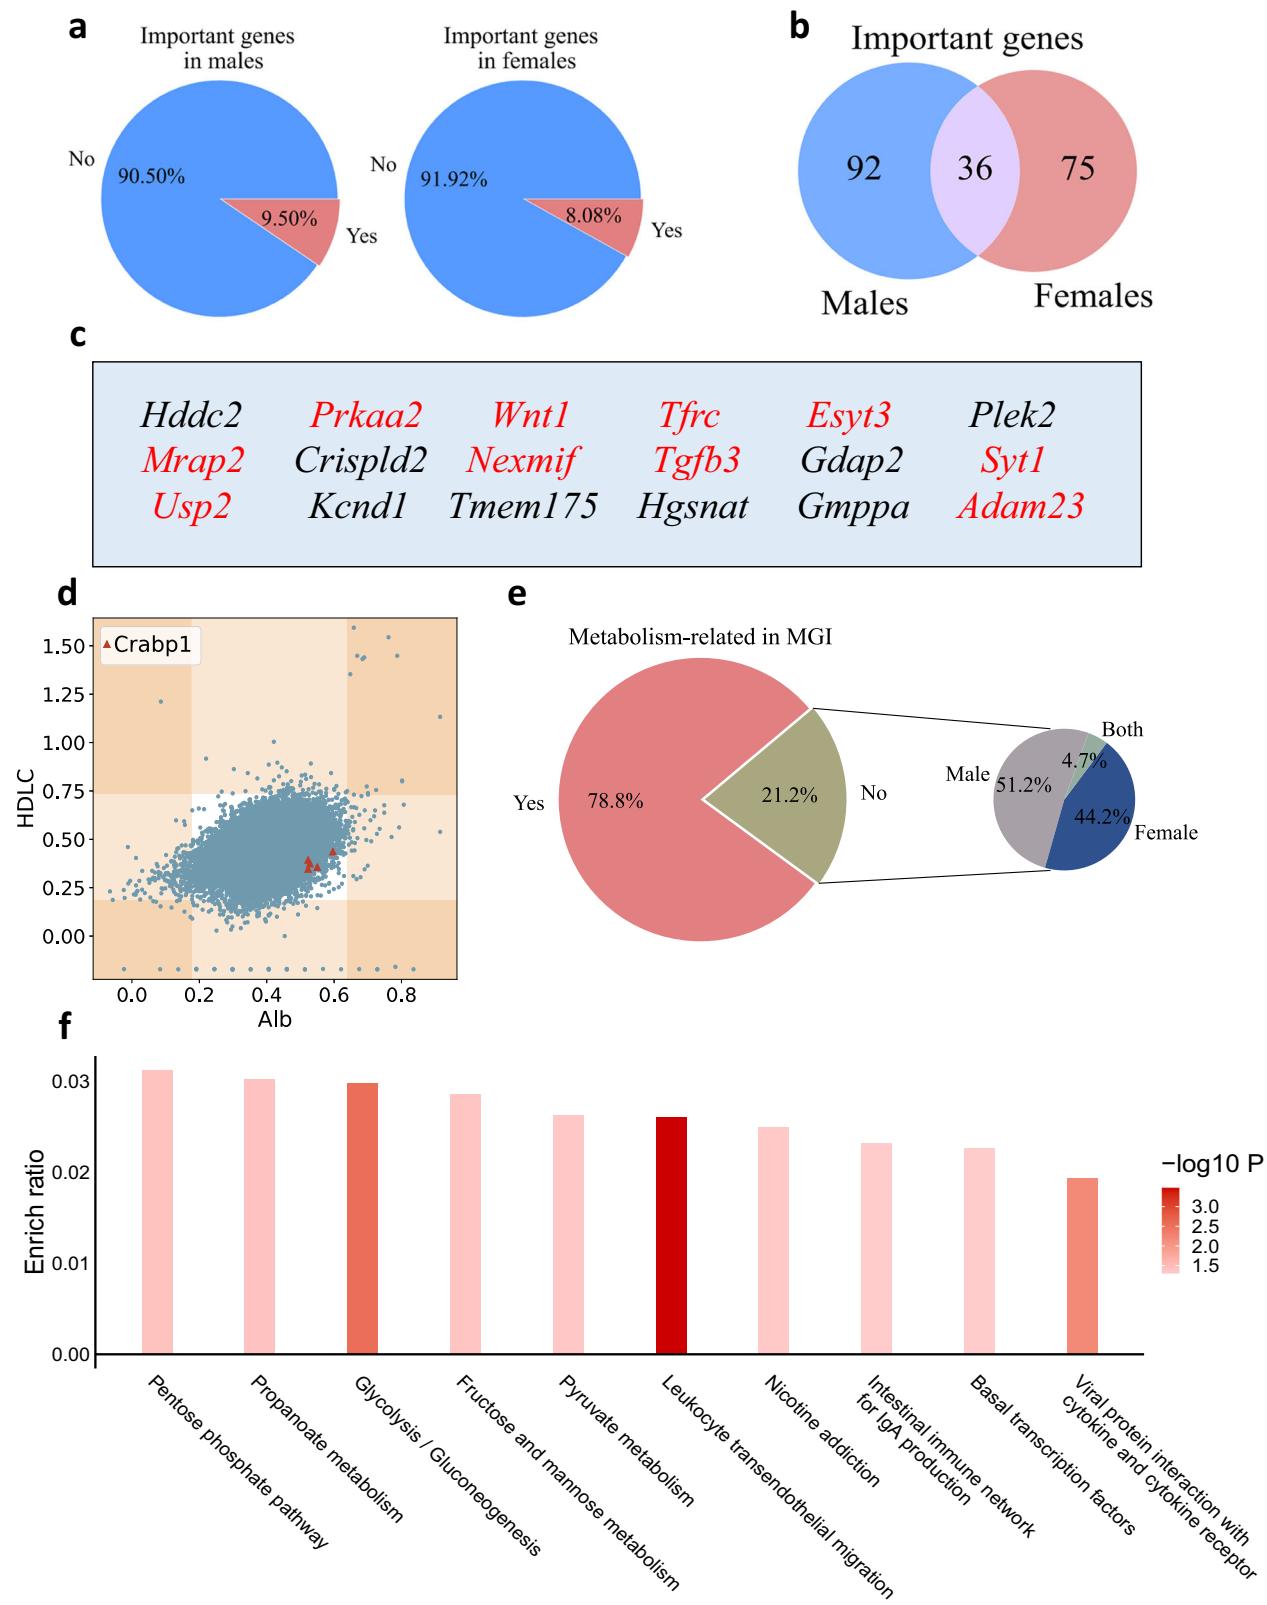


Fig. 5 | ODBAE identifies new metabolic genes from a metabolism-related dataset of knockout mice. a The proportion of important genes identified in the male and female datasets. **b** ODBAE identified 128 important genes in the male dataset and 111 important genes in the female dataset, for a total of 203 important genes. **c** Sexually dimorphic genes identified by ODBAE, with previously reported sexually dimorphic genes highlighted in red. **d** Visualization results of outliers corresponding to gene *Crabc1* in the subspace formed by abnormal parameters. The

absolute z -score values of Alb or HDLC of the data points in the shaded part are ranked in the top 1.2%. **e** Validation results of important genes in the MGI database. 78.8% of the genes are associated with metabolism, and 21.2% are newly identified by ODBAE. Besides, 51.2% and 44.2% new metabolic genes are identified in the male or female datasets, respectively, and 4.7% are detected in both male and female datasets. **f** Histogram of KEGG pathway enrichment for 43 novel genes in mice obtained through KOBAS.

To further explore these correlations, we identified parameter pairs most likely to show joint abnormalities in both male and female mice. In males, 36 different abnormal parameter pairs were observed, with the most common being (Alb, Ph), (Alb, HDLC), and (ALA, HDLC) (Fig. 6a, c). In females, 40 different abnormal parameter pairs were identified, with the most common being (ASA, TP), (Alb, ASA), and (TB, TP) (Fig. 6b, c). Their corresponding knockout genes are shown in Fig. 6d, e. Notably, 22 abnormal parameter pairs were found in both males and females, with the most common pairs being (ALA, ASA), (ALA, HDLC), and (Alb, ASA). These common abnormalities suggest conserved metabolic pathways that are affected by specific gene knockouts in both sexes.

We then investigated the biological significance of the three most frequently perturbed parameter pairs. Previous studies have shown that the ALA/HDLC ratio is associated with the risk of non-alcoholic fatty liver disease (NAFLD) and diabetes^{37–39}. ODBAE identified *Lepr* as one of the genes causing abnormalities in this pair (Fig. 7a). When the ratio of ALA to HDLC was analyzed across all mice, the outliers corresponding to the *Lepr* knockout had a significantly lower ratio compared to the general population (Fig. 7b), consistent with previous findings linking *Lepr* mutations to T2D and NAFLD^{40,41}. KEGG pathway analysis of genes associated with (ALA, HDLC) revealed enrichment in pathways related to NAFLD and metabolic processes, including glycolysis (Fig. 7c and Supplementary Data 3f). For the (Alb, ASA) pair, studies suggest that the ASA/Alb ratio correlates with tumor progression and prognosis in hepatocellular carcinoma (HCC)^{42,43}. ODBAE identified *Ap4m1* as a gene causing abnormalities in this pair (Fig. 7d). The average ratio of ASA to Alb in *Ap4m1* knockout mice was significantly larger than in other mice (Fig. 7e), consistent with the role of *Ap4m1* in HCC⁴⁴. KEGG pathway enrichment further linked (Alb, ASA) disruptions to pathways involved in pyruvate metabolism, glycolysis, and proximal tubule bicarbonate reclamation (Fig. 7f and Supplementary Data 3g). Finally, for genes corresponding to parameter pair (ALA, ASA), the pathway enrichment analysis revealed significant enrichment in Glycosaminoglycan biosynthesis, Starch and sucrose metabolism, and Type I diabetes mellitus (Supplementary Fig. 10a and Supplementary Data 3h), consistent with the previously reported association of parameters ASA to ALA ratio with prediabetes⁴⁵. In addition to these well-characterized parameter pairs, ODBAE also identified novel correlations worth exploring further. For instance, *Taf8* was highlighted as a gene causing abnormalities in the parameter pair (HDLC, TP), though little is known about the specific relationship between these two parameters (Supplementary Fig. 10b). Supplementary Fig. 10c shows that there is a linear relationship between these two parameters, and the knockout of the gene *Taf8* disrupts this relationship. ODBAE's ability to uncover such relationships without prior biological inference makes it a powerful tool for discovering previously unknown metabolic interactions.

Overall, ODBAE systematically identifies and explains complex phenotypes by revealing intrinsic relationships between multiple metabolic parameters, including both linear and non-linear associations. These insights can serve as new indicators for biomedical research and provide a more holistic view of metabolic dysfunction in knockout mice. In addition, the ODBAE framework can be adapted to other datasets to uncover complex phenotypes in different biological systems.

Discussion

In this study, we present ODBAE, a machine learning approach designed to detect and explain complex phenotypes by identifying outliers in high-dimensional biological datasets. The key advantage of ODBAE is its ability to detect two types of outliers: IP, which deviate significantly from expected relationships between variables, and HLP, which are far from the data center. Unlike traditional outlier detection methods that often focus on single abnormal indicators, ODBAE identifies multi-dimensional joint abnormalities, providing a more holistic view of phenotypic disruptions. This approach is particularly powerful when applied to data from knockout mouse models, where phenotypic changes are often subtle or involve complex interactions between physiological parameters.

One of the most significant contributions of ODBAE is its potential to identify unexpected phenotypes that may not be apparent when analyzing individual physiological indicators. Traditional methods of phenotype analysis often focus on abnormalities in a single trait^{3–5}, potentially overlooking the coordinated disruptions across multiple parameters that signal underlying biological imbalances. As our understanding of disease and homeostasis evolves, it is increasingly clear that pathological processes often involve system-wide disturbances rather than isolated dysfunctions⁶. ODBAE leverages these insights by integrating correlations between physiological indicators, even when individual indicators remain within normal ranges. This capability allows researchers and clinicians to better understand the interplay of homeostatic mechanisms and how they are disrupted in disease states. For example, when applied to metabolism-related datasets from IMPC, ODBAE successfully identified novel metabolic genes, including those associated with complex, multi-parameter phenotypes. These phenotypes would have been difficult to detect using conventional approaches that focus on single-parameter abnormalities. Moreover, ODBAE successfully identified sex-dimorphic genes, underscoring its utility in advancing the study of sexual dimorphism. ODBAE also uncovered new parameter pairs that tend to exhibit abnormalities together, providing insights into the intrinsic relationships between metabolic pathways. This type of discovery is essential for advancing our understanding of gene function and the pathophysiology of metabolic diseases. Notably, the utility of ODBAE extends beyond knockout mouse models. It can be broadly applied to diverse tabular biological datasets—including disease cohorts, clinical health examination records, and single-cell datasets—to identify complex phenotypes and subtle, multidimensional anomalies.

Importantly, ODBAE offers not just detection, but also anomaly explanation through the use of kernel-SHAP. This feature allows researchers to pinpoint the exact physiological indicators driving the detected abnormalities, facilitating deeper insights into the biological processes involved. By identifying the most significant reconstruction errors and the features contributing to them, ODBAE enables a clearer interpretation of complex phenotypic data. This is particularly relevant in clinical settings, where understanding the root cause of a phenotype can guide diagnosis and treatment decisions.

Although ODBAE demonstrates considerable potential, there are areas that warrant further exploration. First, the current framework assumes a Gaussian distribution in the training data, which may limit its effectiveness in datasets that do not follow this distribution. As shown in Fig. 4b, ODBAE maintains robust performance under moderate levels of noise in Gaussian-distributed datasets. However, its anomaly detection capability progressively declines with increasing noise intensity. Future work should focus on optimizing the model to handle non-Gaussian data and improving the accuracy of anomaly explanations, particularly in cases involving highly complex or subtle phenotypes. Second, although ODBAE could potentially be extended to other data types, such as time-series or imaging data, the current manuscript focuses exclusively on tabular datasets. The model's performance in analyzing other types of data, such as longitudinal time-series or medical imaging, has not yet been empirically tested. Further exploration and validation in these areas are warranted to assess ODBAE's broader applicability. Third, the performance of ODBAE is influenced by both the number of layers and the dimensionality of each layer, suggesting that its architecture can be tuned to further enhance performance in practical applications. However, determining the optimal model configuration remains a challenging task. Moreover, since ODBAE identifies anomalies based on the reconstruction error of individual data points, the choice of threshold for this error is critical to the accuracy of outlier detection. Further optimization of threshold selection is therefore warranted. Finally, ODBAE's use in identifying phenotypes has primarily focused on metabolic parameters. Its performance in detecting phenotypes related to other physiological systems, such as neurological, cardiovascular, or developmental traits, remains unexplored. Broader testing across a range of physiological categories would demonstrate the model's versatility and help identify any system-specific limitations.

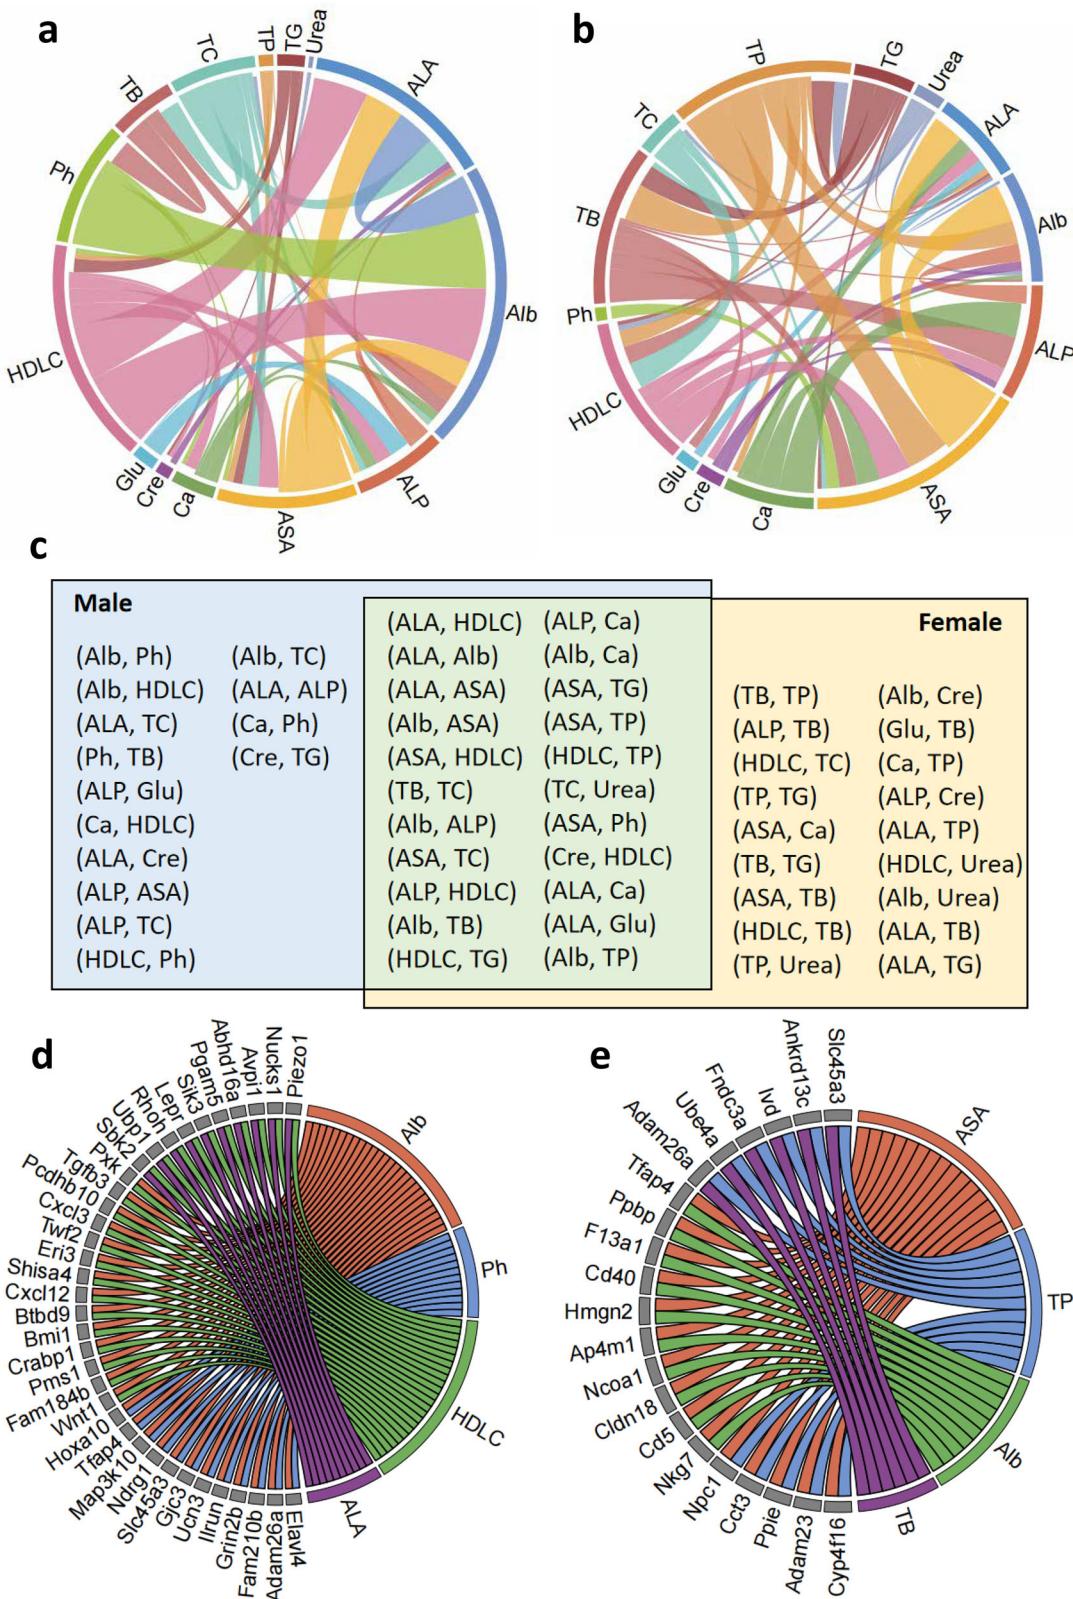


Fig. 6 | ODBAE integrates metabolic parameter pairs tend to simultaneous abnormalities. a, b Chord diagram plotting the inter-connectivity of metabolic parameters for males (a) and females (b). The outer segments represent the metabolic parameters, the size of the arcs connecting the parameters is proportional to the

number of outliers associated. **c** Abnormal parameter pairs observed in males and females. **d, e** Links between genes and simultaneous abnormal metabolic parameters in males (d) and females (e).

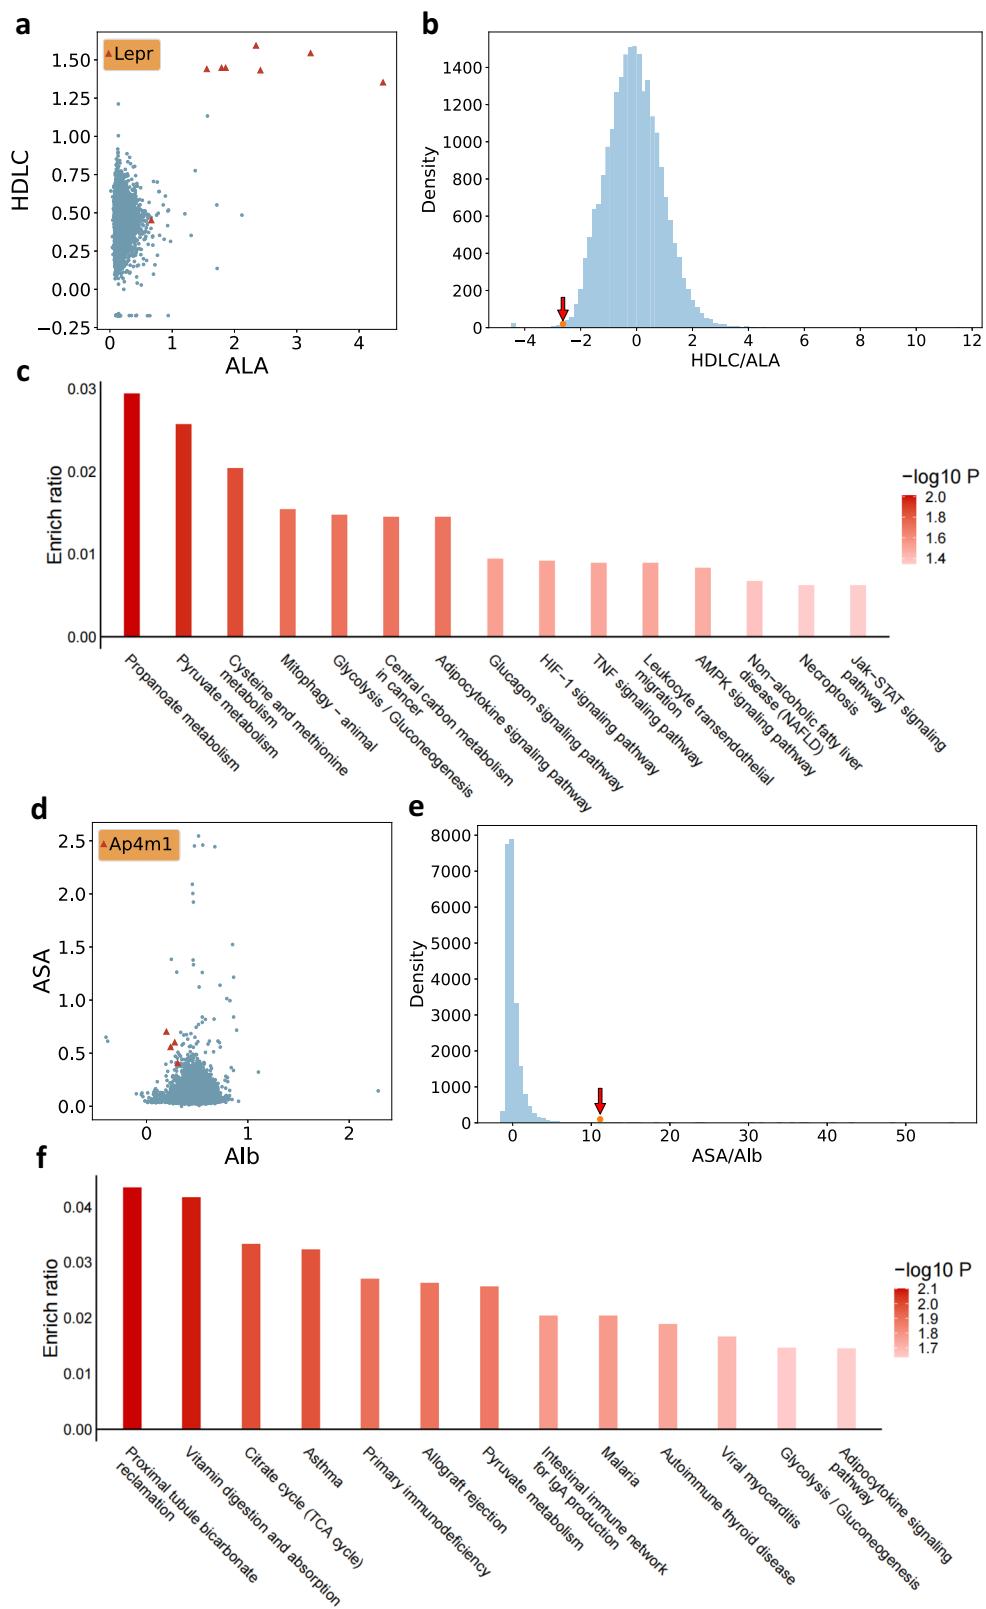


Fig. 7 | Validation of the top three abnormal parameter pairs that occur in both males and females and have the highest overall frequency. a Visualization of outliers corresponding to gene *Lepr* based on abnormal parameter pairs in the male dataset. **b** The distribution of the standardized ratio of ALA to HDLC in mouse metabolism-related dataset, the average ratio of *Lepr* knockout mice is significantly smaller. **c** Histogram of KEGG pathway enrichment for genes corresponding to (ALA, HDLC) in humans. **d** Visualization of outliers corresponding to gene *Ap4m1* based on abnormal parameter pairs in the female dataset. **e** The distribution of the standardized ratio of ASA to Alb in mouse metabolism-related dataset, the average ratio of *Ap4m1* knockout mice is significantly larger. **f** Histogram of KEGG pathway enrichment for genes corresponding to (Alb, ASA) in humans.

(ALA, HDLC) in humans. **d** Visualization of outliers corresponding to gene *Ap4m1* based on abnormal parameter pairs in the female dataset. **e** The distribution of the standardized ratio of ASA to Alb in mouse metabolism-related dataset, the average ratio of *Ap4m1* knockout mice is significantly larger. **f** Histogram of KEGG pathway enrichment for genes corresponding to (Alb, ASA) in humans.

Despite these challenges, ODBAE represents a powerful tool for the analysis of high-dimensional biomedical datasets. Its ability to detect multi-dimensional outliers, explain anomalies, and integrate correlated physiological indicators offers a new way to approach phenotypic screening, particularly in gene knockout models. We believe that ODBAE will facilitate the discovery of previously unrecognized phenotypes, enhancing our understanding of homeostasis and disease processes. By moving beyond the limitations of single-indicator analysis, ODBAE opens up new possibilities for both basic research and clinical diagnostics, allowing for a more comprehensive understanding of complex biological systems.

Methods

Datasets

To demonstrate the comprehensive benefits of ODBAE, we first performed outlier detection on developmental datasets obtained from the IMPC. Subsequently, we assessed ODBAE's ability to identify HLP using 3 synthetic two-dimensional Gaussian-distributed datasets (Fig. 2c–e) and one three-dimensional dataset (Fig. 2i). Finally, we evaluated the model's overall detection accuracy for both IP and HLP using an additional three-dimensional dataset (Supplementary Fig. 5a). For outlier detection on the IMPC developmental dataset, data from wild-type mice were used as the training dataset, while data from knockout mice served as the test dataset. The training dataset comprised 8327 wild-type samples (4126 males and 4201 females), and the test dataset included 26706 knockout samples (13297 males and 13409 females). To evaluate the accuracy of ODBAE in detecting HLP, as well as its overall detection performance for both IP and HLP, all synthetic datasets consisted of 20000 samples each.

To evaluate the advantages of ODBAE over existing methods, we conducted outlier detection using multiple approaches across five low-dimensional synthetic datasets, two high-dimensional synthetic datasets, and two benchmark datasets (the Dry Bean dataset and the Breast Cancer dataset). In the low-dimensional synthetic dataset experiments, we generated three two-dimensional datasets drawn from Gaussian distributions, one from a two-dimensional Laplace distribution, and one from a two-dimensional uniform distribution, each comprising 20000 samples. For the high-dimensional synthetic datasets, we constructed one 50-dimensional and one 100-dimensional Gaussian-distributed dataset, each also consisting of 20000 samples. The Dry Bean dataset contains 13611 samples across seven distinct categories; we randomly selected 2000 samples from a single class as the training dataset, with samples from the remaining six classes designated as outliers. The Breast Cancer dataset comprises 569 samples, categorized as benign or malignant. We randomly selected 300 benign samples as the training data, while malignant samples were treated as outliers.

The metabolism-related dataset used to identify novel metabolic genes and phenotypes included 21234 wild-type mouse samples (10585 males and 10649 females) and 45922 knockout mouse samples (23024 males and 22898 females). Wild-type samples were employed for training, while knockout samples served as the test dataset.

Dataset preprocessing

We used developmental and metabolism-related datasets from IMPC to discover complex phenotypes and identify new genes. The developmental dataset we used was integrated from 8 phenotyping centers (BCM, HMGU, ICS, JAX, MRC Harwell, RBRC, TCP, and UC Davis), while the metabolism-related dataset was from 11 phenotyping centers (BCM, CCP-IMG, HMGU, ICS, KMPC, MARC, MRC Harwell, RBRC, TCP, UC Davis, and WTSI). To eliminate experiment-specific variations, we standardized each phenotyping center for each dataset to have unit variance and zero mean. Then, Min-Max normalization was applied to normalize each physiological parameter within each dataset, mitigating dimensional effects across different physiological parameters. SNPs within a ± 1 kb region of orthologues for the 41 newly identified metabolic genes were downloaded from GWAS Central on September 23, 2023 and filtered based on $-\log(p) \geq 2$.

Definition of HLP and IP

Since the reconstruction of any dataset by the autoencoder can be regarded as a complex regression process, we can divide the outliers in each dataset into HLP and IP⁴⁶. For a certain dataset, we divide the dimension index into two disjoint sets V_1 and V_2 . Without loss of generality, we assume that $V_1 = \{1, 2, \dots, p\}$ and $V_2 = \{p + 1, p + 2, \dots, m\}$. For $i, j \in V_1, i \neq j$, there is no correlation between variables X_i and X_j , while for each $k \in V_2$, there is a map ϕ_k such that variable $X_k = \phi_k(X_1, X_2, \dots, X_p)$. $X_i (i \in V_1)$ is called factor and $X_k (k \in V_2)$ is called response variable. Then we define matrix Q ,

$$Q = \begin{bmatrix} 1 & x_{11} & \dots & x_{1p} \\ 1 & x_{21} & \dots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n1} & \dots & x_{np} \end{bmatrix}.$$

Besides, if we let $Q_i = (1, x_{i1}, \dots, x_{ip})$, the leverage value of the i th sample point is

$$h_i = Q_i(Q^\top Q)^{-1}Q_i^\top, \quad i = 1, 2, \dots, n.$$

If we represent the mean of variable $X_k (k \in V_1)$ as μ_k , and define the mean vector $\tilde{\mu} = (\mu_1, \mu_2, \dots, \mu_p)$, we have

$$Q^\top Q = \begin{bmatrix} 1 & 1 & \dots & 1 \\ x_{11} & x_{21} & \dots & x_{n1} \\ \vdots & \vdots & \ddots & \vdots \\ x_{1p} & x_{2p} & \dots & x_{np} \end{bmatrix} \begin{bmatrix} 1 & x_{11} & \dots & x_{1p} \\ 1 & x_{21} & \dots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n1} & \dots & x_{np} \end{bmatrix} = n \begin{bmatrix} 1 & \tilde{\mu} \\ \tilde{\mu}^\top & C \end{bmatrix}, \quad (1)$$

where $C_{jk} = \frac{1}{n} \sum_{i=1}^n x_{ij}x_{ik}$. Thus, we can obtain the inverse of $Q^\top Q$

$$(Q^\top Q)^{-1} = \frac{1}{n} \begin{bmatrix} 1 + \tilde{\mu}(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top & -\tilde{\mu}(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \\ -(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top & (C - \tilde{\mu}^\top \tilde{\mu})^{-1} \end{bmatrix}. \quad (2)$$

Therefore, if we let $\tilde{x}_i = (x_{i1}, \dots, x_{ip})$, the leverage value of the i th sample point can also be formulated as

$$\begin{aligned} h_i &= \frac{1}{n} Q_i \begin{bmatrix} 1 + \tilde{\mu}(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top & -\tilde{\mu}(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \\ -(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top & (C - \tilde{\mu}^\top \tilde{\mu})^{-1} \end{bmatrix} Q_i^\top \\ &= \frac{1}{n} \left(1 + (\tilde{\mu} - \tilde{x}_i)(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top, (\tilde{x}_i - \tilde{\mu})(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \right) Q_i^\top \\ &= \frac{1}{n} \left(1 + (\tilde{\mu} - \tilde{x}_i)(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{\mu}^\top + (\tilde{x}_i - \tilde{\mu})(C - \tilde{\mu}^\top \tilde{\mu})^{-1} \tilde{x}_i^\top \right) \\ &= \frac{1}{n} + \frac{1}{n} (\tilde{x}_i - \tilde{\mu})(C - \tilde{\mu}^\top \tilde{\mu})^{-1} (\tilde{x}_i - \tilde{\mu})^\top. \end{aligned} \quad (3)$$

If we define $A = C - \tilde{\mu}^\top \tilde{\mu}$, then

$$\begin{aligned} A_{jk} &= \frac{1}{n} \sum_{i=1}^n x_{ij}x_{ik} - \mu_j\mu_k \\ &= \frac{1}{n} \sum_{i=1}^n (x_{ij} - \mu_j)(x_{ik} - \mu_k) \\ &= \frac{n-1}{n} \text{COV}(X_j, X_k), \quad j, k \in V_1. \end{aligned} \quad (4)$$

Furthermore, we have $A = \frac{n-1}{n} \Sigma_f$ and $A^{-1} = \frac{n}{n-1} \Sigma_f^{-1}$, where Σ_f is the covariance matrix of the factor space. Then, based on Eq. (3), we get

$$h_i = \frac{1}{n} + \frac{1}{n-1} (\tilde{x}_i - \tilde{\mu}) \Sigma_f^{-1} (\tilde{x}_i - \tilde{\mu})^\top. \quad (5)$$

Obviously, for each sample point, its leverage value is proportional to its Mahalanobis distance in the factor space. In addition, the Cook distance of the i th sample point based on the k th ($k \in V_2$) dimension is

$$C_{ik} = \frac{(n-p)(x_{ik} - \hat{x}_{ik})^2 h_i}{p \sum_{i=1}^n (x_{ik} - \hat{x}_{ik})^2 (1 - h_i)^2}, \quad i = 1, 2, \dots, n. \quad (6)$$

For each dataset, HLP are sample points with high leverage value and IP refer to sample points with high Cook distance. The leverage value of the i th sample point is also proportional to the Mahalanobis distance of the factor space. Therefore, HLP can also be identified by the Mahalanobis distance in the factor space. It is noteworthy that for certain datasets, like those conforming to a multivariate Gaussian distribution, the set V_2 may be empty. In such instances, all outliers in the dataset are identified as HLP.

Outlier detection of ODBAE

The autoencoder model in ODBAE is formed of an encoder f and a decoder g . If we denote the input data point and the output data point as variable X and \hat{X} , respectively, where $X = (X_1, X_2, \dots, X_m)^\top$ and $\hat{X} = (\hat{X}_1, \hat{X}_2, \dots, \hat{X}_m)^\top$, and the number of sample points is n , then we have $\hat{X} = g \circ f(X)$. Besides, \mathcal{X} represents the whole input dataset and $X \in \mathcal{X}$. x_i represents the i th sample point of the input dataset and its corresponding reconstruction result is \hat{x}_i , $x_i = (x_{i1}, x_{i2}, \dots, x_{im})^\top$ and $\hat{x}_i = (\hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{im})^\top$. Ideally, we would like to train the autoencoder with the training dataset to minimize the training loss function, the generally used training loss function is MSE which can be represented as

$$L_{MSE}(\omega, b) = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{x}_i)^\top (x_i - \hat{x}_i) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m (x_{ij} - \hat{x}_{ij})^2.$$

Where ω represents the weight between the input layer and the output layer and b is the bias value. The purpose of training process is to guarantee that the intrinsic information of the training dataset can be learned and most of the normal sample points can be well reconstructed. It is important to note that MSE-trained autoencoders are insufficient in outlier detection. Through rigorous theoretical analysis, we demonstrated how this limitation can be addressed. The analysis process is based on the following assumption.

Assumption 1. We assume that the input variable X follows m -dimensional Gaussian distribution.

As we know, autoencoders will mainly focus on recovering the principal components of a dataset eventually^{47,48}, it is necessary to use the unit orthogonal eigenvectors as the basis vectors to further study the influence of the loss function on the outlier detection results of the autoencoder. Under Assumption 1, we represent the mean vector of the variable X as $\mu = (\mu_1, \mu_2, \dots, \mu_m)^\top$. If the eigenvalues of the covariance matrix Σ_x are $\lambda_1, \lambda_2, \dots, \lambda_m$, the corresponding unit orthogonal eigenvectors (i.e., principal directions) are $\eta_1, \eta_2, \dots, \eta_m$ and the new coordinate of the variable X is $Y = (Y_1, Y_2, \dots, Y_m)^\top$ when the unit orthogonal eigenvectors are the basis vectors, then we can denote an orthogonal matrix P , where the i th column of P is unit eigenvector η_i . It's obvious that $Y = P^{-1}X = P^\top X$ and $Y_k = \eta_k^\top X$, $k = 1, 2, \dots, m$. Thus, in the new coordinate system, the variable of input data X is converted to variable Y . Let $v = (v_1, v_2, \dots, v_m)^\top$ and Σ_Y represent the mean vector and covariance matrix of variable Y , respectively, we have $v = P^\top \mu$ and $v_k = E(Y_k) = \eta_k^\top \mu$, $k = 1, 2, \dots, m$. Furthermore, we

can gain the variance of the k th element of variable Y is

$$\begin{aligned} D(Y_k) &= \text{COV}(Y_k, Y_k) \\ &= E((Y_k - v_k)(Y_k - v_k)) \\ &= E(\eta_k^\top (X - E(X)) \eta_k^\top (X - E(X))) \\ &= \eta_k^\top E((X - E(X))(X - E(X))^\top) \eta_k \\ &= \eta_k^\top \Sigma_x \eta_k \\ &= \eta_k^\top \lambda_k \eta_k \\ &= \lambda_k, \end{aligned} \quad (7)$$

and the covariance of Y_i and Y_j ($i \neq j$) is

$$\begin{aligned} \text{COV}(Y_i, Y_j) &= E((Y_i - E(Y_i))(Y_j - E(Y_j))) \\ &= E(\eta_i^\top (X - E(X)) \eta_j^\top (X - E(X))) \\ &= \eta_i^\top E((X - E(X))(X - E(X))^\top) \eta_j \\ &= \eta_i^\top \Sigma_x \eta_j \\ &= \eta_i^\top \lambda_j \eta_j \\ &= 0. \end{aligned} \quad (8)$$

Thus, the covariance matrix of variable Y is $\Sigma_Y = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_m)$.

Proposition 1. The covariance matrix of variable Y can be derived as

$$\Sigma_Y = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_m)$$

If we denote $v = (v_1, v_2, \dots, v_m)^\top$ as the mean vector of variable Y , since X follows m -dimensional Gaussian distribution, we can conclude that Y also follows m -dimensional Gaussian distribution, i.e., $Y \sim \mathcal{N}(v, \Sigma_Y)$. What's more, it's easy to get that $Y_k \sim \mathcal{N}(v_k, \lambda_k)$, $k = 1, 2, \dots, m$. Then, we denote \hat{Y} as the output variable of Y and its mean vector and eigenvalues of the covariance matrix are $\hat{v} = (\hat{v}_1, \hat{v}_2, \dots, \hat{v}_m)^\top$ and $\hat{\lambda}_k$, $k = 1, 2, \dots, m$, respectively. As is mentioned above, autoencoders will mainly recover the principal components of any datasets, it's reasonable for us to make the following assumption.

Assumption 2. For each element \hat{Y}_k of the output variable \hat{Y} , its data distribution is the same as the corresponding input element Y_k . Since $Y_k \sim \mathcal{N}(v_k, \lambda_k)$, then $\hat{Y}_k \sim \mathcal{N}(\hat{v}_k, \hat{\lambda}_k)$, $k = 1, 2, \dots, m$. Besides, whether each eigenvalue after reconstruction is close to 0 is the same as its corresponding original eigenvalue.

In the following, we will theoretically analyze the restrictions of MSE in detecting HLP. First, we can get the relationship between the difference between the reconstructed data and its mean and the difference between the original data and its mean in each principal component direction.

As we know, the loss function of MSE can be formulated as

$$\begin{aligned} L_{MSE}(\omega, b) &= \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^\top (y_i - \hat{y}_i) \\ &= E(Y^\top Y - 2Y^\top \hat{Y} + \hat{Y}^\top \hat{Y}) \\ &= E(Y^\top Y) - 2E(Y^\top \hat{Y}) + E(\hat{Y}^\top \hat{Y}) \\ &= \sum_{i=1}^m E(Y_i^2) - 2 \sum_{i=1}^m E(Y_i \hat{Y}_i) + \sum_{i=1}^m E(\hat{Y}_i^2), \end{aligned} \quad (9)$$

where y_i and \hat{y}_i represent the i th sample point of the input and output dataset in the new coordinate system, respectively.

First, if the autoencoder can properly reconstruct the input dataset, we will analyse the mean vector of the output variable \hat{Y} . Based on Assumption

2 and Eq. (9), we have

$$\begin{aligned} L_{MSE}(\omega, b) &= \sum_{i=1}^m \lambda_i + \sum_{i=1}^m \hat{\lambda}_i + \sum_{i=1}^m v_i^2 + \sum_{i=1}^m \hat{v}_i^2 - 2 \sum_{i=1}^m v_i \hat{v}_i \\ &\quad - 2 \sum_{i=1}^m COV(Y_i, \hat{Y}_i) \geq \sum_{i=1}^m (v_i - \hat{v}_i)^2 + \sum_{i=1}^m \left(\sqrt{\lambda_i} - \sqrt{\hat{\lambda}_i} \right)^2. \end{aligned} \quad (10)$$

In order to reach the minima of $L_{MSE}(\omega, b)$, the mean vector of output variable \hat{v} must satisfy $\hat{v}_k = v_k, k = 1, 2, \dots, m$. That is to say, the mean vector of output data is the same as input data.

Next, we will specifically analyze the reconstruction error of each data point. According to Assumptions 1 and 2, the input and output variable satisfy $Y_k \sim \mathcal{N}(v_k, \lambda_k)$ and $\hat{Y}_k \sim \mathcal{N}(v_k, \hat{\lambda}_k), k = 1, 2, \dots, m$. Denote $R(Y) = (R_1(Y), R_2(Y), \dots, R_m(Y))^\top$ ($\hat{R}(Y) = (\hat{R}_1(Y), \hat{R}_2(Y), \dots, \hat{R}_m(Y))^\top$) as the difference between variable Y (\hat{Y}) and its mean vector v , i.e., $R(Y) = Y - v$ and $\hat{R}(Y) = \hat{Y} - v$. It's easy to see that $R_k(Y) \sim \mathcal{N}(0, \lambda_k)$ and $\hat{R}_k(Y) \sim \mathcal{N}(0, \hat{\lambda}_k)$. Then the following equation can be obtained.

$$\begin{aligned} L_{MSE}(\omega, b) &= E((Y - \hat{Y})^\top (Y - \hat{Y})) \\ &= E((R(Y) - \hat{R}(Y))^\top (R(Y) - \hat{R}(Y))) \\ &= E(R^\top(Y)R(Y)) - 2E(R^\top(Y)\hat{R}(Y)) + E(\hat{R}^\top(Y)\hat{R}(Y)) \\ &= \sum_{i=1}^m E(R_i^2(Y)) + \sum_{i=1}^m E(\hat{R}_i^2(Y)) - 2 \sum_{i=1}^m E(R_i(Y))E(\hat{R}_i(Y)) \\ &\quad - 2 \sum_{i=1}^m COV(R_i(Y), \hat{R}_i(Y)) \\ &= \sum_{i=1}^m \lambda_i + \sum_{i=1}^m \hat{\lambda}_i - 2 \sum_{i=1}^m \rho_i \sqrt{\lambda_i} \sqrt{\hat{\lambda}_i}, \end{aligned} \quad (11)$$

where ρ_k represents the correlation coefficient of $R_k(Y)$ and $\hat{R}_k(Y)$, and $-1 \leq \rho_k \leq 1, k = 1, 2, \dots, m$. In Eq. (11), only when $\rho_k = 1$, can $L_{MSE}(\omega, b)$ reach its minima. In this case, there is a positive linear correlation between $R_k(Y)$ and $\hat{R}_k(Y)$, specifically, there exist constants a_k and c_k such that $\hat{R}_k(Y) = a_k R_k(Y) + c_k$, then $E(\hat{R}_k(Y)) = a_k E(R_k(Y)) + c_k$ and $D(\hat{R}_k(Y)) = a_k^2 D(R_k(Y))$. Since $E(\hat{R}_k(Y)) = E(R_k(Y)) = 0$, $D(R_k(Y)) = \lambda_k$ and $D(\hat{R}_k(Y)) = \hat{\lambda}_k$, we can obtain $c_k = 0$ and $a_k = \frac{\sqrt{\hat{\lambda}_k}}{\sqrt{\lambda_k}}$.

Therefore, we can obtain the following proposition.

Proposition 2. Under Assumptions 1 and 2, the difference between the reconstructed data and its mean is proportional to the difference between the original data and its mean in each principal component direction. The specific relationship is

$$\hat{R}_k(Y) = \frac{\sqrt{\hat{\lambda}_k}}{\sqrt{\lambda_k}} R_k(Y). \quad (12)$$

Where $R_k(Y) = Y_k - v_k$ and $\hat{R}_k(Y) = \hat{Y}_k - \hat{v}_k$. Based on Eq. (12), we can further determine the specific reconstruction error of each input data point.

Proposition 3. The reconstruction error of each input data point measured by MSE can be formulated as

$$\begin{aligned} W(Y) &= (Y - \hat{Y})^\top (Y - \hat{Y}) \\ &= \sum_{i=1}^m \left(\sqrt{\lambda_i} - \sqrt{\hat{\lambda}_i} \right)^2 \frac{R_i^2(Y)}{\lambda_i}. \end{aligned} \quad (13)$$

Ideally, we hope each principal component contributes equally to the reconstruction error of any input data point. Recall the definition of $R(Y)$, it's easy to know that $\frac{R_k(Y)}{\lambda_k} \sim \mathcal{N}(0, 1)$, so Eq. (13) indicates that the reconstruction degree of each principal component by the autoencoder will affect the proportion of this principal component in the reconstruction error. This brings a lot of uncertainty for HLP detection.

Proposition 4. HLP are difficult to identify as outliers if the input dataset is completely reconstructed.

When the input dataset is well reconstructed, in Eq. (13), it means that $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} = 0, k = 1, 2, \dots, m$. Then, the reconstruction error of each data point in training dataset is 0. As the result, each data point in the training dataset is considered normal. However, HLP exist in any dataset follows m -dimensional Gaussian distribution. Generally, the larger the Mahalanobis distance is, the more abnormal the data point is. Therefore, although completely reconstructing the training dataset is beneficial for IP detection, it ignores the detection of HLP.

Proposition 5. HLP whose values in each dimension are within the normal range can not be identified as outliers. Besides, most HLP will be detected in the direction corresponding to the worst-recovered principal component, but in the direction of the well-recovered principal components, the anomalies are often ignored.

From Eq. (13), we know $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k}$ is equivalent to the weight of the reconstruction error of the k th principal component direction of each data point to the total reconstruction error. Therefore, if the data reconstruction in the i th principal component direction is better than that in the j th principal component direction, then $\sqrt{\lambda_i} - \sqrt{\hat{\lambda}_i} < \sqrt{\lambda_j} - \sqrt{\hat{\lambda}_j}, i \neq j$. As the result, the detected outliers are more distributed in the j th principal component direction.

Proposition 6. If we ensure the differences between the eigenvalues of the covariance matrix of the original dataset and their corresponding reconstructed results in the direction of each principal component are equal, the value of the reconstruction error for each data point will be proportional to its Mahalanobis distance.

If we suppress complete reconstruction of the autoencoder, then $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} > 0, k = 1, 2, \dots, m$. Based on Eq. (12), we can obtain

$$(Y_k - \hat{Y}_k)^2 = \frac{\left(\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} \right)^2}{\lambda_k} (Y_k - v_k)^2. \quad (14)$$

It means that for each principal component direction of the input data, there is a reconstruction error. Specifically, the further away the values are from the mean, the larger the reconstruction error. Therefore, in each principal component direction, the outliers detected by the reconstruction error are the same as the outliers defined by the Gaussian distributed data. Besides, if

$\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} = \beta > 0, k = 1, 2, \dots, m$, according to Eq. (13), the

reconstruction error of each input data point is

$$W(Y) = (Y - \hat{Y})^\top (Y - \hat{Y}) = \beta^2 \sum_{i=1}^m \frac{R_i^2(Y)}{\lambda_i}. \quad (15)$$

Since $\Sigma_Y = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_m)$, it's easy to obtain $\Sigma_Y^{-1} = \text{diag}(\frac{1}{\lambda_1}, \frac{1}{\lambda_2}, \dots, \frac{1}{\lambda_m})$. So the Mahalanobis distance of the input variable Y is

$$M(Y) = (Y - \nu)^\top \Sigma_Y^{-1} (Y - \nu) = \sum_{i=1}^m \frac{R_i^2(Y)}{\lambda_i}. \quad (16)$$

It indicates that the reconstruction error of each data point is proportional to its Mahalanobis distance. Therefore, if we control $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} = \beta > 0$, $k = 1, 2, \dots, m$, almost all HLP will be converted to IP and cannot be well reconstructed. As the result, the detected HLP will evenly distributed in each principal component direction.

Based on the above discussion, we will propose a new loss function that adds an appropriate penalty term based on MSE to balance the reconstruction of the autoencoder.

In fact, if the intrinsic dimension of the dataset is l , then there will be l eigenvalues that are not close to 0. We denote the l eigenvalues that are not close to 0 as λ'_k and their corresponding reconstruction results are $\hat{\lambda}'_k$, $k = 1, 2, \dots, l$. According to Assumption 2, $\hat{\lambda}'_k$ is also not close to 0. Then, we consider two losses in our training loss function. One of them is $L_{MSE}(\omega, b)$, which aims to reconstruct the input dataset well. In addition, we define the other loss $L_{EIG}(\omega, b) = \sum_{i=1}^l (\sqrt{\lambda'_i} - \sqrt{\hat{\lambda}'_i} - \beta)^2$ which can avoid the autoencoder from completely reconstructing the input dataset. $\beta > 0$ is a hyperparameter that can adjust the degree of data reconstruction. The final training loss function is a combination of the two:

$$L(\omega, b) = \theta_1 L_{MSE}(\omega, b) + \theta_2 L_{EIG}(\omega, b). \quad (17)$$

Here, $\theta_1, \theta_2 > 0$ are hyperparameters that need to be predetermined. In practice, desirable results are usually obtained if we set $\theta_1 = 0.008$ and $\theta_2 = 1$.

Specification of outlier ratio

ODBAE identifies outliers based on reconstruction error, with larger reconstruction errors indicating greater abnormality. In the implementation of ODBAE, we provide two strategies for determining the reconstruction error threshold. Given a dataset of N samples, when a user-defined outlier proportion δ is specified, the model classifies the top $N \times \delta$ samples with the highest reconstruction errors as outliers. Alternatively, when the threshold is determined automatically, the model assesses the distribution of reconstruction errors. If the errors approximately follow a Gaussian distribution, outliers are identified based on the 3σ rule. Otherwise, kernel density estimation is employed to derive the probability density function of the reconstruction errors, and data points with reconstruction errors exceeding the 95th percentile are considered anomalous.

Anomaly explanation of ODBAE

For each outlier detected by ODBAE, it's necessary for us to find its abnormal parameters. For ODBAE, it obtained abnormal parameters based on the highest reconstruction errors and SHAP values. Specifically, for each outlier, we first sorted its parameters in descending order based on the reconstruction errors. Then, if the sum of the reconstruction errors of the first n parameters was greater than 50% of the total reconstruction error for the outlier, these n parameters were considered potentially anomalous. If $n = 1$, we used SHAP values to identify additional parameter that have the greatest impact on this anomalous parameter. If the SHAP value of these two parameters exceeded the set threshold, they would be considered abnormal parameters; otherwise, only the parameter with the highest reconstruction error was considered an abnormal parameter. However, if $n > 1$, we would use SHAP values to obtain additional parameter that have the greatest

impact on the parameter with the highest reconstruction error. If the SHAP value of these two parameters exceed the set threshold, they would be considered abnormal parameters; otherwise, the top 2 parameters with the highest reconstruction error were considered anomalous parameters. In our work, the threshold for SHAP values was set to the mean of the SHAP values during the process of anomaly explanation for all outliers.

Identification of outliers using z-score method

In developmental and metabolism-related datasets, we first identified mice with anomalies using the z-score method. For example, there were 14 metabolic parameters in metabolism-related datasets. Then, the z-score value of each mouse for each metabolic parameter was calculated as $Z_{ij} = \frac{X_{ij} - \mu_i}{\sigma_i}$, where Z_{ij} represented the z-score of the j th mouse for the i th metabolic parameter, X_{ij} was the value of the i th metabolic parameter for the j th mouse, μ_i and σ_i were the mean and standard deviation of the i th metabolic parameter, respectively. The larger the absolute value of Z_{ij} , the more the i th metabolic parameter value of the j th mouse deviates from the mean. Therefore, for each metabolic parameter, mice with large absolute z-score values were labeled as outliers.

Identification of mice with low BMI values

We calculated the BMI values for all knockout mice, and the mean BMI of the abnormal mice after *Ckb* knockout in female mice was lower than that of 97.14% of the mice. Therefore, there were still a small number of mice with lower BMIs from a total of 264 single-gene knockout mouse strains, and only 13 (4.92%) single-gene knockout mouse strains had more than 50% of mice with BMI values lower than the mean BMI of gene *Ckb* knockout mice. Among these 13 genes, 7 genes were identified as important genes by ODBAE, and for the remaining 6 genes, a portion of mice from each corresponding gene knockout strain were detected as outliers.

Validation of sexually dimorphic genes

When applied to the IMPC metabolic dataset, ODBAE identified a total of 203 metabolism-associated genes, with a marked divergence observed between male- and female-specific genes (Fig. 5b). To further validate the biological relevance of these findings, we first screened for genes exhibiting significant sexual dimorphism. In our framework, a gene was considered significant if the proportion of outliers among its corresponding knockout strains exceeded 50%. To ensure that the observed abnormalities were indeed attributable to the gene and not random fluctuations, we initially restricted our analysis to genes with sufficient sample sizes in both sexes. Given that the majority of knockout strains in both male and female datasets comprised seven individuals (Supplementary Fig. 11a, b), we retained genes with at least seven samples per sex. Subsequently, to identify genes with pronounced sexual dimorphism, we focused on genes for which the outlier ratio exceeded 50% in only one sex and was zero in the other. This yielded a final set of 18 genes showing strong evidence of sex-specific metabolic effects.

To further evaluate whether the sexual dimorphism observed in these 18 genes had been previously reported, we conducted a literature search in PubMed using each gene name in combination with the keywords "sexual dimorphism", "sexual", "sex", "male", or "female". This search revealed that 10 of the 18 genes had been previously documented as exhibiting sexual dimorphism.

Comprehensive performance evaluation

We compared ODBAE with MSE-AE, MAE-AE, and DAGMM for the detection results of IP and HLP, and computed their AUC and AP scores by regarding the IP and HLP as positive.

We implemented DAGMM with minimal modification so that it adapt to our datasets. Besides, for DAGMM, we kept its number of layers and the dimensions of each layer consistent with ODBAE, MSE-AE, and MAE-AE.

For each experiment, the structure of the autoencoder was as follows. The dimension of the hidden layer was the intrinsic dimension of the input dataset, and its activation function was ReLU. Besides, the activation

function of the output layer was sigmoid. All datasets were normalized before outlier detection. During the autoencoder training process, Adam⁴⁹ was used to optimize the loss function and the learning rate was 10^{-3} .

To evaluate the performance of various outlier detection methods across a range of predefined outlier ratios, each method was executed 10 times per ratio. For each run, we recorded the AUC and AP scores. The final comparison was based on the mean AUC and AP scores averaged across the ten replicates for each method.

Selection of hyperparameters θ_1, θ_2 in training loss function

Our new training loss function is formulated as $L(\omega, b) = \theta_1 L_{MSE}(\omega, b) + \theta_2 L_{EIG}(\omega, b)$. It can be seen that our loss function contains two loss terms. The function of $L_{MSE}(\omega, b)$ is to drive the autoencoder to reconstruct the input dataset as well as possible, while $L_{EIG}(\omega, b)$ aims at suppressing complete reconstruction of the dataset by the autoencoder and balance the degree of reconstruction in the direction of each principal component of the input dataset. Obviously, when the value of $\theta_1: \theta_2$ is large, the autoencoder can reconstruct the input dataset perfectly and will adversely affect the detection of HLP. However, if the value of $\theta_1: \theta_2$ is very small, the autoencoder will have a poor reconstruction effect on the input dataset, which will bring disadvantages to IP detection. Therefore, it is very necessary for us to ensure that the value of $\theta_1: \theta_2$ is reasonable, so that the autoencoder can have a good detection effect for HLP and IP at the same time.

In our work, we found that when $\theta_1: \theta_2 = 0.008: 1$, most experiments could obtain the desired results. In order to study the sensitivity of this ratio, we changed its base to see how different base affects the accuracy of outlier detection. To be specific, if the value of base was c , we set the values of θ_1 and θ_2 to $0.008c$ and c , respectively. We conducted experiments on the Dry Bean dataset. For a fixed base c , we assessed the outlier detection performance of ODBAE across different outlier ratios δ ($\delta \in \{0.05, 0.1, 0.15, 0.2, 0.25\}$), and subsequently computed the average score across all ratios. The results show that the AUC and AP scores obtained by ODBAE varied only marginally with changes in the base value (Supplementary Table 1), indicating that θ_1 and θ_2 are insensitive to the change of the base.

Determination of hyperparameter β

Although hyperparameter $\beta > 0$ in Eq. (17) is beneficial for HLP detection, as the value of β increases, the data reconstruction ability of the autoencoder will become worse and worse, which is similar to the model not being well fitted during regression analysis. This adversely affects the detection of IP. Therefore, in the following, it's important for us to determine how to choose the appropriate value of β .

If the intrinsic and actual dimensions of the dataset are equal, i.e., $l = m$, then most of the outliers in the dataset are HLP. In this case, considering that $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} > 0$ and $\hat{\lambda}_k > 0$, $k = 1, 2, \dots, m$, the value of the hyperparameter β should satisfy $0 < \beta < \min_{1 \leq i \leq m}(\sqrt{\lambda_i})$. According to the previous analysis, as long as $\beta > 0$, the detection effect of HLP can be improved. Meanwhile, in order to maintain the detection effect of IP, the value of β must be small enough.

According to Eq. (12), if $\beta = 0$, then $\hat{R}_k(Y) = R_k(Y)$, i.e., $\hat{Y}_k = Y_k$, $k = 1, 2, \dots, m$. However, we usually add nonlinear activation functions to the autoencoder to learn nonlinear features in the input dataset. For example, in our work, we added sigmoid activation function to the output layer. Since commonly used nonlinear activation functions are saturated, when the input values are very large or very small, the corresponding output values hardly change with the change of the input values. As the result, for each dimension of the input data, when the values are very large or very small, their reconstruction results are poor. Otherwise, the values can be reconstructed well. Then, most HLP detected by the autoencoder are anomalous in one dimension, but HLP whose values in each dimension are within the normal range are ignored. It can be seen from this that when β is small enough, it not only adversely affects the detection effect of HLP, but also rarely improves the detection effect of IP. That is to say, there exists

$\xi > 0$, when $\beta < \xi$, there is not much improvement in the detection effect of IP.

If we determine the value of ξ and let $\beta = \xi$, on the one hand, the detection effect of the autoencoder for IP can be maintained. On the other hand, we can balance the reconstruction of the autoencoder and avoid the influence of the saturation of the activation function on the data reconstruction at the same time, which can improve the detection effect for HLP.

In our work, the structure of the autoencoder and the setting of hyperparameters are described above. For each dimension of any dataset, input values between 0.15 and 0.85 are hardly affected by the saturation of the nonlinear activation function. Therefore, we hope that the value of β selected can make the output value of each dimension between 0.15 and 0.85. As the result, for each dimension, the interval length of the output value is 0.7 times the interval length of the input value. Before outlier detection, the input dataset is normalized, so the input value of each dimension is between 0 and 1. Recall that $Y_k \sim \mathcal{N}(v_k, \lambda_k)$ and $\hat{Y}_k \sim \mathcal{N}(v_k, \hat{\lambda}_k)$, $k = 1, 2, \dots, m$, the probability that Y_k is distributed in $(v_k - 3\sqrt{\lambda_k}, v_k + 3\sqrt{\lambda_k})$ is 0.9974, so its distribution interval length is approximately $6\sqrt{\lambda_k}$. Similarly, the distribution interval length of \hat{Y}_k is approximately $6\sqrt{\hat{\lambda}_k}$. So we can control $\sqrt{\hat{\lambda}_k} = 0.7\sqrt{\lambda_k}$, then $\sqrt{\lambda_k} - \sqrt{\hat{\lambda}_k} = 0.3\sqrt{\lambda_k}$, $k = 1, 2, \dots, m$.

Based on the above discussion, we can summarize how to determine the value of the hyperparameter β when $l = m$. If $\max_{1 \leq i \leq m}(0.3\sqrt{\lambda_i}) \leq \min_{1 \leq i \leq m}(\sqrt{\lambda_i})$, we can set $\beta = \max_{1 \leq i \leq m}(0.3\sqrt{\lambda_i})$. Otherwise, we set $\beta = \min_{1 \leq i \leq m}(\sqrt{\lambda_i})$.

In practical applications, we usually encounter this type of dataset whose intrinsic dimension l is smaller than the actual dimension m . In this case, there are only l eigenvalues of the covariance matrix of the dataset that are not close to 0, we denote them as λ'_k , $k = 1, 2, \dots, l$. In the training process, we only need to control such l eigenvalues and make $\sqrt{\lambda'_k} - \sqrt{\hat{\lambda}'_k} = 0.3\sqrt{\lambda'_k}$, $k = 1, 2, \dots, l$. Besides, since $\lambda'_k > 0$, β should satisfy $0 \leq \beta \leq \min_{1 \leq i \leq l}(\sqrt{\lambda'_i})$. Actually, if $l < m$, there will be some data points that have a large impact on the reconstruction ability of the autoencoder (i.e., IP), it is very important to not only improve the detection effect of HLP, but also maintain the detection effect of IP. Therefore, we have to make sure that the value of β is small enough but not equal to 0. For convenience, we also set $0 < \beta < \min_{1 \leq i \leq m}(\sqrt{\lambda_i})$ in this case. Finally, the determination scheme for the value of β is the same as when $l = m$.

Statistics and reproducibility

All data analysis were performed using Python 3.7.7 and R 3.4.3 software. To ensure the reproducibility of our findings, we provide detailed descriptions of the data generation process, including sample sizes used in the synthetic data experiments. For experiments involving real-world datasets, we report the sources of the data, the preprocessing procedures, and all relevant analysis details. When comparing the outlier detection performance of ODBAE against other models, we ran each model 10 times under fixed outlier ratios and recorded the AUC scores from each run. The final performance comparisons are based on the average AUC scores across runs to ensure robustness and reproducibility. Similarly, to assess the impact of dataset dimensionality, noise level, and model architecture on ODBAE's performance, we report the average AUC scores obtained from 10 independent runs under each condition.

Ethics statement

The gene knockout mouse data used in this study were obtained from the International Mouse Phenotyping Consortium (IMPC), and no experiments were conducted here. All IMPC research centers make every effort to minimize animal suffering through careful housing and husbandry, with detailed information available at the IMPC portal: <https://www.mousephenotype.org/about-impc/arrive-guidelines>. Ethics statements from individual research centers are provided in Supplementary Data 4.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The raw data from IMPC was downloaded from https://ftp.ebi.ac.uk/pub/databases/impc/all-data-releases/release-13.0/cores/statistical-raw-data_20201217_163937.tar. The Dry Bean Dataset was downloaded from <https://archive.ics.uci.edu/ml/datasets/Dry+Bean+Dataset>. The Breast Cancer Dataset was downloaded from <https://www.kaggle.com/datasets/uciml/breast-cancer-wisconsin-data>. Additionally, all source data are provided in the Supplementary Data or are available at <https://github.com/YafeiShen2022/ODBAE>.

Code availability

The implementation of ODBAE, along with the main analysis code, is available at <https://github.com/YafeiShen2022/ODBAE>.

Received: 11 October 2024; Accepted: 29 August 2025;

Published online: 02 October 2025

References

- Kurbatova, N., Mason, J. C., Morgan, H., Meehan, T. F. & Karp, N. A. Phenstat: a tool kit for standardized analysis of high throughput phenotypic data. *PLoS ONE* **10**, e0131274 (2015).
- Deans, A. R. et al. Finding our way through phenotypes. *PLoS Biol.* **13**, e1002033 (2015).
- Rozman, J. et al. Identification of genetic elements in metabolism by high-throughput mouse phenotyping. *Nat. Commun.* **9**, 288 (2018).
- Swan, A. L. et al. Mouse mutant phenotyping at scale reveals novel genes controlling bone mineral density. *PLoS Genet.* **16**, e1009190 (2020).
- Meehan, T. et al. Disease model discovery from 3,328 gene knockouts by the International Mouse Phenotyping Consortium. *Nat. Genet.* **49**, 1231–1238 (2017).
- Nicholson, G. et al. Multivariate phenotype analysis enables genome-wide inference of mammalian gene function. *PLoS Biol.* **20**, e3001723 (2022).
- Banfalvi, G. Homeostasis. In *Homeostasis - tumor - metastasis*. 1–92 (Springer, Dordrecht, 2014).
- Elgart, M. E. M. et al. Correlations between complex human phenotypes vary by genetic background, gender, and environment. *Cell Rep. Med.* **3**, 100844 (2022).
- Bulik-Sullivan, B. et al. An atlas of genetic correlations across human diseases and traits. *Nat. Genet.* **47**, 1236–1241 (2015).
- Kanai, M. et al. Genetic analysis of quantitative traits in the Japanese population links cell types to complex human diseases. *Nat. Genet.* **50**, 390–400 (2018).
- Lyu, R. et al. Geslm algorithm for detecting causal SNPs in GWAS with multiple phenotypes. *Brief. Bioinforma.* **22**, bbab276 (2021).
- Sattarov, T., Herurkar, D. & Hees, J. Explaining anomalies using denoising autoencoders for financial tabular data. Preprint at <https://arxiv.org/abs/2209.10658> (2022).
- Lu, W. et al. Unsupervised sequential outlier detection with deep architectures. *IEEE Trans. Image Process.* **26**, 4321–4330 (2017).
- Wei, Y. et al. Lstm-autoencoder-based anomaly detection for indoor air quality time-series data. *IEEE Sens. J.* **23**, 3787–3800 (2022).
- Ding, K., Li, J., Bhanushali, R. & Liu, H. Deep anomaly detection on attributed networks. In *19th SIAM International Conference on Data Mining*, 594–602 (Society for Industrial and Applied Mathematics, 2019).
- Chen, J., Milot, L., Cheung, H. M. C. & Martel, A. L. Unsupervised clustering of quantitative imaging phenotypes using autoencoder and gaussian mixture model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 575–582 (Springer International Publishing, 2019).
- Xu, D., Ricci, E., Yan, Y., Song, J. & Sebe, N. Learning deep representations of appearance and motion for anomalous event detection. In *British Machine Vision Conference*, (eds Xie, X., Jones, M. W. & Tam, G. K. L.) 8.1–8.12 (BMVA Press, 2015).
- Le, V.-T. & Kim, Y.-G. Attention-based residual autoencoder for video anomaly detection. *Appl. Intell.* **53**, 3240–3254 (2022).
- Zhou, C., Jia, Y. & Motani, M. Optimizing autoencoders for learning deep representations from health data. *IEEE J. Biomed. Health Inform.* **23**, 103–111 (2019).
- Liu, Y. et al. Integration of human protein sequence and protein-protein interaction data by graph autoencoder to identify novel protein-abnormal phenotype associations. *Cells* **11**, 2485 (2022).
- Imon, A. H. M. R. & Hadi, A. S. Identification of multiple high leverage points in logistic regression. *J. Appl. Stat.* **40**, 2601–2616 (2013).
- Antwarg, L., Miller, R. M., Shapira, B. & Rokach, L. Explaining anomalies detected by autoencoders using shapley additive explanations. *Expert Syst. Appl.* **186**, 115736 (2021).
- Shyu, M., Chen, S., Sarinnapakorn, K. & Chang, L. A novel anomaly detection scheme based on principal component classifier. In *Proceedings of the IEEE Foundations and New Directions of Data Mining Workshop*, 172–179 (IEEE, 2003).
- Shalabi, L. A. A., Shaaban, Z. & Kasasbeh, B. Data mining: a preprocessing engine. *J. Comput. Sci.* **2**, 735–739 (2006).
- Karp, N. A. et al. Prevalence of sexual dimorphism in mammalian phenotypic traits. *Nat. Commun.* **8**, 15475 (2017).
- Novelli, E. L. B. et al. Anthropometrical parameters and markers of obesity in rats. *Lab. Anim.* **41**, 111–119 (2007).
- Hu, Y., Rosa, G. J. M. & Gianola, D. A gwas assessment of the contribution of genomic imprinting to the variation of body mass index in mice. *BMC Genomics* **16**, 576 (2015).
- Macêdo, A. P. A. et al. Murinometric measurements and retroperitoneal adipose tissue in young rats exposed to the high-fat diet: Is there correlation? *Braz. J. Biol.* **81**, 246–250 (2021).
- Rahbani, J. F. et al. Creatine kinase b controls futile creatine cycling in thermogenic fat. *Nature* **590**, 480–485 (2021).
- Zong, B. et al. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International Conference on Learning Representations* (2018).
- Beck, T., Rowlands, T., Shorter, T. & Brookes, A. J. Gwas central: an expanding resource for finding and visualising genotype and phenotype data from genome-wide association studies. *Nucleic Acids Res.* **51**, D986–D993 (2023).
- Morris, A. P. et al. Large-scale association analysis provides insights into the genetic architecture and pathophysiology of type 2 diabetes. *Nat. Genet.* **44**, 981–990 (2012).
- Ried, J. S. et al. A principal component meta-analysis on multiple anthropometric traits identifies novel loci for body shape. *Nat. Commun.* **7**, 13357 (2016).
- Willer, C. J. et al. Discovery and refinement of loci associated with lipid levels. *Nat. Genet.* **45**, 1274–1283 (2013).
- Cotsapas, C. et al. Pervasive sharing of genetic effects in autoimmune disease. *PLoS Genet.* **7**, e1002254 (2011).
- Bu, D. et al. Kobas-ii: intelligent prioritization and exploratory visualization of biological functions for gene enrichment analysis. *Nucleic Acids Res.* **49**, W317–W325 (2021).
- Cao, C. et al. The nonlinear correlation between alanine aminotransferase to high-density lipoprotein cholesterol ratio and the risk of diabetes: a historical Japanese cohort study. *BMC Endocr. Disord.* **23**, 124 (2023).
- He, S. et al. Alanine aminotransferase to high-density lipoprotein cholesterol ratio is positively correlated with the occurrence of diabetes in the Chinese population: a population-based cohort study. *Front. Endocrinol.* **14**, 1266692 (2023).

39. Qiu, J. et al. The newly proposed alanine aminotransferase to high-density lipoprotein cholesterol ratio has shown effectiveness in identifying non-alcoholic fatty liver disease. *Front. Endocrinol.* **14**, 1239398 (2023).
40. Yang, M. et al. Variations in the obesity gene “lepr” contribute to risk of type 2 diabetes mellitus: evidence from a meta-analysis. *J. Diab. Res.* **2016**, 5412084 (2016).
41. Li, X. et al. Gene polymorphisms associated with non-alcoholic fatty liver disease and coronary artery disease: a concise review. *Lipids Health Dis.* **15**, 53 (2016).
42. Wu, Y. et al. Prognostic nomograms for hbv-related bclc 0-a stage hepatocellular carcinoma incorporating aspartate aminotransferase to albumin ratio. *Scand. J. Gastroenterol.* **58**, 789–797 (2023).
43. Peng, W. et al. Preoperative aspartate aminotransferase to albumin ratio correlates with tumor characteristics and predicts outcome of hepatocellular carcinoma patients after curative hepatectomy: a multicenter study. *BMC Surg.* **22**, 307 (2022).
44. Peng, Y., Li, X., Kang, K. & Zhou, Y. Ap4m1 as a prognostic biomarker associated with cell proliferation, migration and immune regulation in hepatocellular carcinoma. *Cancer Cell Int.* **23**, 235 (2023).
45. Cao, C. et al. Nonlinear relationship between aspartate aminotransferase to alanine aminotransferase ratio and the risk of prediabetes: a retrospective study based on Chinese adults. *Front. Endocrinol.* **13**, 1041616 (2022).
46. Chatterjee, S. & Hadi, A. S. Influential observations, high leverage points, and outliers in linear regression. *Stat. Sci.* **1**, 379–393 (1986).
47. Bao, X., Lucas, J., Sachdeva, S. & Grosse, R. B. Regularized linear autoencoders recover the principal components, eventually. In *Neural Information Processing Systems*. (Larochelle, H., Ranzato, M. and Hadsell, R., Balcan, M. F. & Lin, H.) 6971–6981 (Curran Associates Inc., 2020).
48. Oftadeh, R., Shen, J., Wang, Z. & Shell, D. Eliminating the invariance on the loss landscape of linear autoencoders. In *International Conference on Machine Learning*. {eds. III, Hal Daumé and Singh, A.} 7405–7413 (PMLR, 2020).
49. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, {ICLR} 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. (eds Bengio, Y. & LeCun, Y.) (2015).

Acknowledgements

This work was supported by the National Natural Science Foundation of China (12471467), the National Key Research and Development Program of China (2018YFA0801103) to Ling Yang, Interdisciplinary basic Frontier Innovation Program of Suzhou Medical College of Soochow University (YXY2303022), National Center for International Research (2017B01012) to Ying Xu, the National Natural Science Foundation of China (32100931), Jiangsu Province Engineering Research Center of Development and Translation of Key Technologies for Chronic Disease Prevention and Control to Tao Zhang. The authors would like to thank Professor Moli Huang's lab

members (School of Biology and Basic Medical Sciences, Soochow University), Professor Yu Tang (School of Mathematical Sciences, Soochow University) and Professor Huanfei Ma (School of Mathematical Sciences, Soochow University) for their valuable suggestions and kind help.

Author contributions

Y.S., Y.X., and L.Y. designed the research. Y.S. performed framework implementation and generated results. Y.S., T.Z., Z.L., and K.K. analyzed the data and interpreted the results. L.Y. and Y.X. provided overall supervision and direction of the work. Manuscript writing was done by Y.S. and T.Z.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42003-025-08817-y>.

Correspondence and requests for materials should be addressed to Ying Xu or Ling Yang.

Peer review information *Communications Biology* thanks Enrico Capobianco, Georgios N. Dimitrakopoulos and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editors: Michiaki Hamada and Aylin Bircan, Jasmine Pan.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025