



OPEN
DATA DESCRIPTOR

A self-monitoring analysis and reporting technology dataset of 147,496 hard disks

Shuting Wei¹, Hongzhang Yang¹✉, Zhengguang Chen² & Ping Wang³

In order to study hard disk failure prediction, this paper introduces SMART-Z, a dataset comprising 147,496 pieces of hard disk SMART data periodically collected by a large distributed video data center in China in the enterprise application environment from March 2017 to February 2018. There are 65 types of hard disk models, including 712 failure disks and the rest are healthy disks. To minimize business interference, data acquisition utilized predefined peak-hour exclusion lists, multi-dimensional monitoring, and an intelligent fuse strategy to effectively guarantee the stable operation. Compared to similar open source datasets, SMART-Z additionally discloses the critical value, worst value, device IP, business scenario, drive letter name and other attributes, which is helpful for researchers to track the change of hard disk capacity through time series analysis, and realize regional equipment distribution statistics by business scenario dimensions, thereby building hard disk failure prediction model. After verification, our dataset exhibits only 5.3% blank data, outperforming the 2022 Backblaze ST4000DM000 hard disk data, where the blank value accounts for 14.78% of the total data.

Background & Summary

Hard disk failures are generally divided into two types: predictable and unpredictable. The latter may occur occasionally and is difficult to prevent, such as sudden chip failure, mechanical impact, etc. However, the wear of motor bearings and the aging of disk magnetic media are all predictable situations, and these abnormal phenomena can be detected several days or even weeks in advance. Self Monitoring Analysis and Reporting Technology (SMART)¹ uses a threshold method to monitor faults, which compares the current indicators (or properties) of the hard drive with a set threshold. If the threshold is exceeded, the hard drive will issue an alert to the operating system, giving the user some time to transfer important data to other storage devices. In 1995, Compaq submitted this technical solution to the Small Form Factor (SFF) committee for standardization, becoming a technical standard for automatically monitoring the integrity of hard disk drives and reporting potential issues.

The technical principle of SMART is to collect information from various sensors on the hard drive and store the information in the system service area of the hard drive. This area is generally located in the first few dozen physical tracks of the hard drive's 0 physical surface, and is written into the relevant internal management program by the hard drive manufacturer. In addition to the SMART information table, this also includes low-level formatting programs, encryption and decryption programs, self-monitoring programs, automatic repair programs, etc. The monitoring software used by the user reads SMART information through a command called "SMART Return Status" and does not allow any user to modify the information. The ID code for SMART detection of hard drives represents various detection parameters of the hard drive in two hexadecimal digits.

Although SMART threshold can alert hard disk failures, its accuracy is insufficient. In recent years, researchers have proposed many methods for predicting hard disk failures based on SMART data and machine learning algorithms. The hard disk failure prediction dataset they use is usually derived from the SMART data of hard disks published by Backblaze². Taking the data released in the fourth quarter of 2024 as an example, the storage system has 30 different brands and models of hard drives, and this dataset has become a widely recognized third-party dataset for hard drive failure prediction. The Backblaze dataset contains the raw value RAW VALUE for each SMART ID, the regularized value VALUE, and a label indicating whether a fault has occurred. Although the number and scale of hard drives are large, the category only includes SATA hard drives, and the critical values and worst values for each attribute have not been publicly disclosed. The SMART-Z category disclosed in

¹Tianjin University of Technology, Tianjin, China. ²ZTE Corporation, Shenzhen, China. ³Peking University, Beijing, China. ✉e-mail: yanghongzhang@email.tjut.edu.cn

this article includes SATA hard drives and SAS hard drives, with a total of 65 different brands and models of hard drives. The types of hard drives are relatively diverse and superior to the Backblaze dataset.

Similarly, in 2013, Baidu released a SMART dataset containing 23395 SATA hard drives³, all of which were Seagate ST31000524N models. The collection frequency of this dataset is once every hour, with only SMART data collected from the hard drive and no other indicators. The disadvantage is that all the hard drives are of the same brand and model, and the prediction model proposed by it is not applicable to other brand models.

Sidi Lu *et al.*⁴ first demonstrated that by combining hard disk location information with SMART attributes, it is possible to highly accurately predict hard disk failures. Because hard drives with similar physical spaces are more likely to be affected by the same environmental factors, such as relative humidity and temperature, and may experience similar levels of vibration. Although vibration is not a part of SMART attributes, it can affect the reliability of the hard drive. Therefore, adding location information can capture hard drives running in similar environments or operating conditions, which may encounter similar fault characteristics.

In addition, the power on duration of the hard drive may also have a certain impact on its lifespan. Long term power on of the hard drive (even when idle) may cause charge leakage and make it prone to overheating, accelerating the aging of the hard drive. Long term operation of the hard drive may also increase the risk of firmware errors.

Therefore, in order to provide more possibilities for hard disk failure prediction methods, SMART-Z has released hard disk SMART data from a large distributed video data center in China from March 2017 to February 2018. The data consists of three.csv files: log_hdsmart_base.csv, log_hdsmart_param.csv and offline_hdsmart.csv, which display the basic information of the hard disk, the SMART information of the hard disk, and the information of the faulty disk. The SMART-Z dataset adds various attributes such as Threshold (Threshold), Worst value (Worst), Device IP (omcip), Business scenario (bureau), disk name (czname), powerontime of offline hard disk (powerontime) based on the Backblaze dataset. Researchers can explore faulty hard drives from multiple dimensions and discover their potential patterns based on the SMART-Z dataset proposed in this article: tracking the trend of hard drive capacity changes through time series analysis; Utilize business scenario dimensions to achieve regional device distribution statistics; Compare the reliability of manufacturer equipment based on the model and type fields; The impact of cumulative power on the faulty hard drive; A hard disk fault prediction model can be constructed based on this to identify high-risk devices in advance, and the mapping relationship between their drive letter names and device IP provides a critical path for fault location.

Due to the exclusivity of obtaining SMART hard disk instructions (such as smartctl), calling them can cause system performance jitter. Collecting SMART on all hard disks at a fixed cycle will inevitably consume a large amount of system resources in a short period of time. After testing, the delay time of collecting SMART of 10000 hard drives simultaneously is almost unacceptable, during which the system can hardly respond to any normal read and write operations and occasionally experiences downtime. Considering the above issues of obtaining SMART for hard drives, this article isolates adjacent data collection based on time windows during collection, and tries to disperse the performance jitter caused by collecting hard drive information within the time window. While meeting the required collection frequency for hard drive fault prediction, the negative impact of collecting hard drive information on front-end business is minimized as much as possible.

Methods

In this section, we describe the collection process behind the hard disk fault prediction dataset. Through an automated hard disk monitoring system, we periodically collect SMART data and its associated management information (such as business scenarios), and process it to obtain the hard disk fault prediction dataset presented in this article. The dataset disclosed in this paper comes from ZTE Corporation, the third author's organization. We periodically collected 147,496 SMART data of hard disks in the enterprise application environment from March 2017 to February 2018.

Data acquisition. The system adopts a sampling strategy based on time windows and a fixed period collection method, with each hard drive collecting data once a day. In order to avoid performance interference during peak periods of video playback services, this article sets high activity periods (11:00–14:00, 18:00–23:00) through a predefined peak period exclusion list (BUSY_TIME_LIST), and uses a real-time clock verification module to avoid high concurrency periods and pause data collection.

The data collection program runs directly at the system level through scripts, first automatically identifying all active block storage devices in the system. The device discovery process relies on native command-line tools, which can dynamically adapt to various interface types including SATA and SAS, ensuring comprehensive coverage of the acquisition range and adaptation to different hard disk devices, with good hardware compatibility. After completing device identification, the system enters a periodic polling process. Within the valid sampling window, the system sets a 20 second timeout threshold for a single hard disk operation to prevent the process from hanging. Considering that backup methods represented by snapshots often trigger timers at the hour or half hour, collection during this period should be avoided. Therefore, each collection time is 5 minutes to 25 minutes and 35 minutes to 55 minutes. Considering the concurrency of data collection, it is advisable to avoid collecting all hard drives simultaneously. Therefore, at intervals of 10 seconds, SMART information of 50–100 hard drives should be collected each time, while device mounting information (such as device IP) at the operating system layer should be read. Therefore, the overall collection process is to collect smart information of 50–100 hard disks every 10 seconds at 5–25 minutes and 35–55 minutes per hour in non peak time, so that each hard disk can be scanned once a day at different time periods instead of once every hour. For business scenario fields, the system will perform real-time comparison with the CMDB configuration library to ensure the accuracy of logical attribution relationships. Considering business downturns, triggering time-consuming operations such as updating models, merging records, and deleting samples outside of peak periods. SMART data is

collected through the smartctl tool, covering indicators highly related to health status such as remapping sector count, number of sectors to be processed, and hard disk temperature. It is worth noting that this article includes three types of hard drives: SATA_HDD, SATA SSD, and SAS_HDD, which have slightly difference when collecting SMART attributes. Both SATA_HDD and SATA SSD types of hard drives can directly view all SMART information of the hard drives through the collection statement smartctl -a /dev/<device>, but the content of the collection results varies greatly. The detailed SMART feature differences can be seen in Table 3. SAS_HDD disks are usually managed through LSI/Avago/Broadcom RAID controllers, which require confirmation of the disk number through lsblk -g in advance, and then use the collection statement smartctl -a -d megaraid, /dev/ to get SMART information for the hard disk.

Critical value and worst value. In addition to recording the original indicators, this method also designs a dynamic extreme value tracking mechanism to capture the maximum abnormal values that have occurred in each ID item during the entire monitoring period when the hard disk is running. The worst value is the maximum abnormal value that has occurred for each ID item during the operation of the hard disk. By performing real-time comparison operations on the peak statistics of data degradation during hard disk operation, the value is continuously refreshed. The worst value can characterize the extreme pressure or degradation states experienced by the device, which often cannot be obtained solely through a single snapshot observation. Usually, the worst value is equal to the current value. If there is a significant fluctuation in the worst value, it indicates that the hard drive has experienced errors or harsh working conditions (such as temperature). In addition, the critical value is a threshold value specified by the hard disk manufacturer to represent the reliability of a certain project, also known as a threshold, which is calculated through a specific formula. If the current value of a parameter approaches the critical value, it means that the hard drive will become unreliable, which may lead to data loss or hard drive failure.

Outlier detection. This article effectively ensures the stable operation of the hard disk data acquisition system through multi-dimensional monitoring and intelligent circuit breaker strategies. In terms of zombie process detection, the system innovatively constructed a three-dimensional state monitoring model, which comprehensively judges from three dimensions: process lifecycle, resource occupancy pattern, and I/O blocking characteristics. When the running time of a process exceeds twice the collection cycle, the CPU usage rate remains above 90% for five cycles, or the I/O waiting queue exceeds 10 tasks, the system will automatically mark it as an abnormal process. In response to these abnormal processes, the system's abnormal processes perform forced termination operations and update the list of abnormal hard disks in real time to ensure that the problematic hard disks can be isolated and processed in a timely manner. In terms of timeout protection, a basic timeout threshold of 20 seconds is used, and logarithmic level dynamic adjustments are made based on the current number of managed hard drives. Through this intelligent timeout management, the system can maintain stable data collection capabilities in complex and changing operating environments, without interrupting normal operations due to strict threshold settings or allowing abnormal processes to occupy resources due to loose thresholds.

At the same time, in timeout or abnormal process scenarios, the system cannot read SMART data or can only read some unreliable old SMART cache data, so the system clearly cannot read valid SMART data. In this case, the script prioritizes stability and would rather skip suspicious hard drives than record unreliable data. The existing logic avoids script freezing or data pollution caused by hard disk offline, which meets the robustness requirements of monitoring scripts.

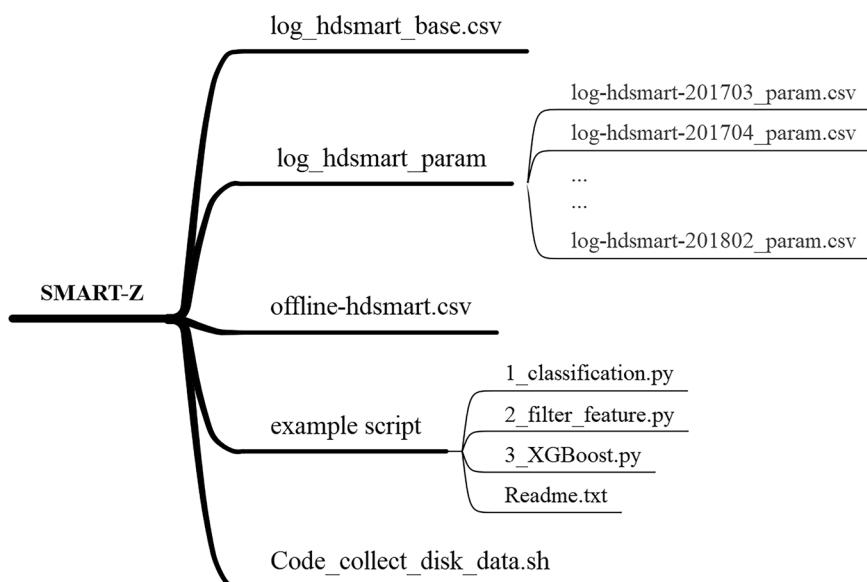
Key parameter configuration. The main parameter configurations of the script during the collection process are shown in Table 1.

Data Records

This dataset⁵ can be found in the open science framework (<https://doi.org/10.17605/OSF.IO/24Y6G>). The DOI is DOI 10.17605/OSF.IO/24Y6G, and assign the license CC-By Attribution 4.0 International. It includes hard disk data processed by a large distributed video data center in China from March 2017 to February 2018. The dataset is composed of multiple files, each corresponding to the specific data categories listed below. The detailed file organization structure is shown in Fig. 1.

- log_hdsmart_base.csv: This file contains 14,525,830 basic information about the hard disks. Among them, each row of data is a scan acquisition of a hard disk on a certain day, including other information of the hard disk except SMART attribute. See log_hdsmart_param.csv for the specific SMART of the hard disk. Main fields: hard disk serial number (keyHDid), collection time (keyScanTime), device IP (omcip), drive letter name (czname), hard disk type (Type), hard disk model (Model), hard disk capacity (Capacity), business scenario (bureau).
- log_hdsmart_param.csv: This file contains 195,840,912 SMART information for hard drives in different months. Among them, each SMART_ID is a separate line, and several consecutive lines may be the same SMART acquisition from the same disk. Because SMART_IDs are different, several lines are generated. Main fields: hard disk serial number (keyHDid), collection time (keyScanTime), SMART_ID, SMART_NAME, standard value corresponding to the current SMART_ID (value), worst value (worst), critical value (threshold), raw value corresponding to the current SMART_ID (rawvalue), SMART warning level (type), SMART update method (update).
- offline-hdsmart.csv: This file contains 712 offline hard drive information, with main fields including brand, business scenario (bureau), hard drive capacity (capacity), device type (devicetype), disk type (disktype), hard drive serial number (hdid), hard drive model (hdnum), the disk offline time (lastlogtime), device IP (omcip), and power on time (powerontime).

| Parameter Category | Configuration Item | Typical Values |
|--|-------------------------|----------------|
| Timeout | timeout_time | 20 s |
| Sleep time between collecting hard disk data | collect_sleep_time | 10 s |
| Maximum number of error log files | error_log_number_max | 30 |
| Upload file failed maximum number of retries | retry_times_max | 10 |
| Upload file retry maximum sleep time | upload_retry_sleep_time | 60 s |

Table 1. Key parameter configuration.**Fig. 1** The actual organization of dataset in detail.

- example script: This folder contains the simple script file and its description document for processing CSV provided by us.
 - 1_classification.py: This script extracts three different types of hard drives and saves them in different CSV files to facilitate researchers in predicting failures of different types of hard drives.
 - 2_filter_feature.py: This script performs feature filtering on CSV files containing SMART attributes, mainly deleting features in the dataset that are all null and have a variance of 0.
 - 3_XGBoost.py: The script first marks the faulty disk as 1 and the healthy disk as 0, and performs algorithm level sample imbalance processing on the dataset, using XGBoost for fault prediction.
- Code_collect_disk_data.sh: The shell script is the code to collect the hard disk failure prediction data set, and can be run in the Linux system environment.

Due to intermittent disconnection caused by loose power supply or connection wires in the `log_hdsmart_base.csv` file, the initial time of hard disk acquisition for each serial number may be different. The same disk may appear multiple times in the file due to daily collection. As time passes, the same hard drive business scenario and device IP may change.

In the `log_hdsmart_param.csv` file, `type` is used to distinguish the fault warning level of SMART parameters, with two values: `Old_age` and `Pre fail`. `Old_age` represents that the SMART attribute reflects the regular aging or long-term wear and tear of the hard drive. Changes in such attributes are expected normal phenomena and will not cause immediate failure of the hard drive, but abnormal values may indicate that the lifespan is approaching its end. And `Pre fail` represents that the SMART attribute is an early warning indicator for hard disk failures, and abnormal values may indicate an imminent failure. That is to say, if the SMART attribute marked as `Pre fail` is abnormal, it is usually judged as a fault, while if the SMART attribute marked as `Old_age` is abnormal, it requires long-term observation to make a judgment. It should be noted that the same SMART attribute of different models of hard drives from the same manufacturer may have different warning levels. For example, the SMART ID 7 of WDC WD4002FYYZ-01B7CB0 is marked as `Old_age`, and the SMART ID 7 of WDC WD4002FYYZ-01B7CB0 is marked as `Pre_fail`.

It should be emphasized that in the `offline-hdsmart.csv` file, all hard drives are offline hard drives. If the serial number of a hard drive does not appear in the file, it is assumed to be a healthy hard drive. Given the complexity of hard disk failures, there is no universally accepted definition for hard disk failures⁶. Therefore, in

| (a)SAS_HDD | | | |
|-------------|---------------------|--------|-------|
| Brand | Model | health | fault |
| HGST | HUS724040ALS640 | 3769 | 6 |
| | HUS724020ALS640 | 9854 | 26 |
| | HUS726040AL5210 | 15486 | 89 |
| | HUS726060AL5210 | 5869 | 12 |
| | HUS726020AL5210 | 14109 | 12 |
| | HUH728080AL5200 | 120 | 0 |
| SEAGATE | ST4000NM0023 | 12642 | 42 |
| | ST300MM0006 | 372 | 1 |
| | ST1000NM0023 | 1377 | 5 |
| | ST2000NM0023 | 5 | 0 |
| | ST4000NM0025 | 174 | 0 |
| TOSHIBA | AL13SEB300 | 408 | 4 |
| WD | WD2001FYYG-01SL3 | 943 | 6 |
| | WD4001FYYG-01SL3 | 2826 | 36 |
| | WD4001FYYG-79SL3 | 20 | 1 |
| | WD1001FYYG-01SL3 | 3 | 0 |
| (b)SATA_SSD | | | |
| INTEL | SSDSC2BX200G4 | 629 | 47 |
| | SSDSC2BB480G4 | 1841 | 18 |
| | SSDSC2BB480G6 | 2409 | 9 |
| | SSDSC2BB240G4 | 207 | 0 |
| | SSDSC2BB300G4 | 2 | 0 |
| | SSDSC2BB800G7 | 282 | 0 |
| | SSDSC2BX400G4 | 30 | 0 |
| | SSDSC2BB480G7O | 12 | 0 |
| TOSHIBA | THNSN8200PCSE | 69 | 0 |
| SanDisk | SDLF1DAR-480G-1HA1 | 1880 | 0 |
| | SDLF1DAM800G-1HHS | 220 | 0 |
| (c)SATA_HDD | | | |
| HGST | HTE545032A7E680 | 486 | 4 |
| | HTE545032A7E380 | 873 | 11 |
| | HUS724020ALS640 | 0 | 0 |
| | HUS726020ALE610 | 3247 | 8 |
| | HUS726040ALE610 | 4769 | 25 |
| | ST2000NM0011 | 251 | 0 |
| | ST9500620NS | 32 | 0 |
| | ST2000NM0055-1V4104 | 263 | 0 |
| | ST4000NM0035-1V4107 | 71 | 0 |
| | ST6000NM0115-1YZ110 | 576 | 0 |
| | ST4000NM0033-9ZM170 | 4806 | 0 |
| | ST2000NM0033-9ZM175 | 3709 | 0 |

Continued

in this article, we only declare all offline hard drives. The offline standard of the hard disk in this paper refers to the frequent occurrence of one or more types of hard disk: mechanical noise, high temperature, serious abnormality of SMART properties of the hard disk, data loss, read-write failure, slow read-write (that is, the waiting time of read-write requests in the queue exceeds 1.5 seconds for 20 consecutive times), frequent offline, file system log printing Medium_Error, I/O_Error, Critical_target_error, Metadata_I/O_Error and other alarm information, unable to mount, no response to read-write operations, etc., of which slow read-write is the most common.

Some researchers believe that the failure of the hard disk refers to the problem of the hard disk itself, such as the damage of the magnetic head and the non rotation of the motor. Others believe that in addition to the failure of the hard disk itself, it also includes the failure caused by the external causes of the hard disk, such as the drop of the line caused by the poor contact of the hard disk cable interface. Therefore, we did not directly define and set the fault tag, which was analyzed by the researchers themselves according to the hard disk fault requirements and combined with the log_hdsmart_base.csv and log_hdsmart_param.csv files. For the convenience of the following description, we will temporarily assume that all offline disks are faulty disks.

This dataset includes three types of hard drives: SAS_HDD, SATA_HDD, and SATA_SSD, totaling 65 hard drive models. The models and their corresponding numbers of healthy and faulty disks are shown in Table 2.

| (a)SAS_HDD | | | |
|------------|--------------------|--------|-------|
| Brand | Model | health | fault |
| WDC | WD2000FYYZ-01UL1B2 | 9221 | 44 |
| | WD2003FYY5-02W0B1 | 859 | 12 |
| | WD4002FYYZ-01B7CB0 | 7738 | 33 |
| | WD4000FYYZ-01UL1B2 | 13523 | 125 |
| | WD4000FYYZ-01UL1B0 | 436 | 4 |
| | WD4000FYYZ-01UL1B1 | 2401 | 17 |
| | WD2000FYYZ-01UL1B1 | 1220 | 13 |
| | WD4000FYYZ-01UL1B3 | 2583 | 4 |
| | WD2004FBYZ-01YCBB0 | 1433 | 12 |
| | WD2000FYYZ-03UL1B3 | 12 | 0 |
| | WD2000FYYZ-01UL1B3 | 3 | 0 |
| | WD2004FBYZ-01YCBB1 | 8774 | 58 |
| | WD4000FYYZ-03UL1B3 | 68 | 4 |
| | WD2000FYYZ-01UL1B0 | 919 | 5 |
| | WD2003FYY5-02W0B0 | 8 | 0 |
| | WD2500BHTZ-04JCPV1 | 16 | 0 |
| | WD4001FFSX-68JNUN0 | 3 | 0 |
| | WD1003FBYZ-010FB0 | 2244 | 16 |
| | WD1003FBYX-01Y7B1 | 81 | 0 |
| | WD2002FYPS-01U1B1 | 1 | 0 |
| | WD2005FBYZ-01YCBB1 | 96 | 2 |
| | WD2005FBYZ-01YCBB2 | 31 | 1 |
| | WD2002FYPS-02W3B1 | 1 | 0 |
| | WD2003FYPS-27Y2B0 | 1 | 0 |
| TOSHIBA | MG04ACA600E | 465 | 0 |
| Hitachi | HTE543232A7A384 | 6 | 0 |

Table 2. Hard disk details.

In SATA_HDD type hard drives, there are mainly 18 SMART features, in SAS_HDD type hard drives, there are mainly 4 SMART features, and in SSD type hard drives, there are mainly 26 SMART features. The detailed SMART attributes are shown in Table 3.

From the table data and collection results, there are several points that need to be noted:

- (1) The low-level collection tool SAS basically cannot collect any information, which results in fewer SMART attributes of SAS_HDD.
- (2) Many SMART attributes collected by the low-level collection tool SDD are Unknown Attributes, such as SMART_ID 170-174, 225-227, 234, 243.
- (3) The same ID for different brands may have different meanings. SMART_ID 175 represents Power_Loss_Cap_Test in SEAGATE and Program_Fail_Count_Chip in WDC.
- (4) The same ID of the same brand but different models may have different meanings. For example, in INTEL SSDSC2BB480G6, it represents Runtime_Bad_Block, and in INTEL SSDSC2BB480G4, it represents SATA_Downshift_Count.

It is worth noting that SMART is a technology for analyzing and monitoring the status of various components of a hard drive, such as magnetic heads, motors, disc, etc. However, not all SMART attributes are related to hard drive failures, and some interference features can have a negative impact on prediction accuracy. In order to help researchers eliminate some of them, we analyzed some of them, and the results are shown in Fig. 2.

From Figure (a), it can be seen that there is a clear distinction in the numerical variation of SMART ID 175 between faulty and healthy hard drives, while in Figure (b), the numerical variation of SMART ID 5 is 0 and constant on both types of hard drives, and has no positive effect on distinguishing between faulty and healthy drives. From this, we conclude that the presence of interference features can have an impact on prediction accuracy. In addition, if inappropriate features are used for machine learning modeling, it not only fails to provide positive assistance for fault prediction, but may also cause confusion in the prediction results. Therefore, it is necessary to perform reasonable feature screening on the collected SMART information. On the one hand, redundant attributes can be removed, input dimensions can be reduced, and model complexity can be lowered. On the other hand, it can improve the predictive performance of the model.

In the log_hdsmart_base.csv file, there are a total of 44 business scenario information, which includes all the business scenarios involved in the hard disk in this dataset, as shown in Table 4.

Table 4 shows the information of 44 business scenarios. The specific fields are composed of province_operator_equipment number, in which the operator includes DX, YD and LT. For example, Anhui_DX_10012 refers to that the hard disk is from the data center with the number of 10012 deployed by China DX in Anhui Province.

| Type | ID | SMART attribute | Detailed explanation |
|----------|-----|-------------------------|--|
| SATA_SSD | 5 | Reallocated_Sector_Ct | The number of sectors that have been remapped |
| | 9 | Power_On_Hours | The time when the hard drive is powered on |
| | 12 | Power_Cycle_Count | The number of times the hard drive is powered on/off |
| | 170 | Available_Reservd_Space | Remaining reserved space |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 171 | Program_Fail_Count | The number of failed blocks in Flash programming |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 172 | Erase_Fail_Count | The number of failed blocks to be erased |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 174 | Unsafe_Shutdown_Count | The number of unexpected power outages on the hard drive |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 175 | Power_Loss_Cap_Test | Power outage protection test |
| | | Program_Fail_Count_Chip | The number of failures that occurred during the data writing process |
| | 183 | SATA_Downshift_Count | The number of times the SATA interface rate has decreased due to errors |
| | | Runtime_Bad_Block | The number of newly added bad blocks in SSD during use |
| | 184 | End-to-End_Error | The number of errors in mapping between logical addresses and physical addresses |
| | 187 | Reported_Uncorrect | Errors that cannot be corrected by hardware ECC |
| | 190 | Temperature_Case | Temperature during hard drive operation |
| | | Airflow_Temperature_Cel | Airflow temperature on the surface of the internal disk of the hard drive |
| | 192 | Unsafe_Shutdown_Count | The number of unexpected power outages on the hard drive |
| | | Power-Off_Retract_Count | The number of times the magnetic head returns |
| | 194 | Temperature_Internal | The current temperature inside the hard drive |
| | | Temperature_Celsius | The current temperature inside the hard drive |
| | 197 | Current_Pending_Sector | Number of sectors waiting to be mapped |
| | 199 | CRC_Error_Count | CRC error count |
| | | UDMA_CRC_Error_Count | CRC error count in Ultra DMA mode |
| | 225 | Host_Writes_32MiB | The amount of data written by the host |
| | | Unknown_SSD_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 226 | Workld_Media_Wear_Indic | The average wear level of current particles |
| | | Unknown_SSD_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 227 | Workld_Host_Reads_Perc | The percentage of host read operations to the total workload of the hard drive |
| | | Unknown_SSD_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 228 | Workload_Minutes | Accumulated time of hard disk in workload state |
| | | Power-off_Retract_Count | The number of times the magnetic head returns |
| | 232 | Available_Reservd_Space | Remaining reserved space |
| | 233 | Media_Wearout_Indicator | Medium wear index |
| | 234 | Thermal_Throttle | Hard disk performance throttling |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| | 241 | Host_Writes_32MiB | The amount of data written by the host |
| | | Total_LBAs_Written | Accumulated LBA write count |
| | 242 | Host_Reads_32MiB | The amount of data read by the host |
| | | Total_LBAs_Read | Accumulated LBA reads |
| | 243 | NAND_Writes_32MiB | Total Write of NAND Flash Memory |
| | | Unknown_Attribute | Vendor defined parameters, specific definitions not disclosed |
| SAS_HDD | 9 | Power_On_Hours | The time when the hard drive is powered on |
| | 900 | Grown_Defect_Elements | The number of defect elements dynamically increasing during the use of hard drives |
| | 901 | Read_Errors | The number of errors that occur while reading data |
| | 902 | Write_Errors | The number of errors that occur when writing data |

Continued

The business scenario information contains the physical deployment location information of the hard disk. The hard disk in the same business scenario will be affected by the same environmental factors such as temperature, humidity, dust and particles, which makes the hard disk in the same business scenario can also fail for the same reason once one of the hard disks fails, that is, the hard disk in the same business scenario may form a

| Type | ID | SMART attribute | Detailed explanation |
|----------|-----|-------------------------|--|
| SATA_HDD | 1 | Raw_Read_Error_Rate | Error occurred when the magnetic head reads data from the surface of the disk |
| | 3 | Spin_Up_Time | The time it takes for the spindle motor to start and reach its rated speed |
| | 4 | Start_Stop_Count | The number of times the hard disk spindle motor starts/stops |
| | 5 | Reallocated_Sector_Ct | The number of sectors that have been remapped |
| | 7 | Seek_Error_Rate | Error rate during magnetic head seek |
| | 9 | Power_On_Hours | The time when the hard drive is powered on |
| | 10 | Spin_Retry_Count | Counting of spindle motor attempts to restart |
| | 11 | Calibration_Retry_Count | Magnetic head calibration retry count |
| | 12 | Power_Cycle_Count | The number of times the hard drive is powered on/off |
| | 183 | Runtime_Bad_Block | The number of bad blocks detected during use |
| | 192 | Power-Off_Retract_Count | Accumulated the number of times the magnetic head returns |
| | 193 | Load_Cycle_Count | Magnetic head loading count |
| | 194 | Temperature_Celsius | The current temperature inside the hard drive |
| | 196 | Reallocated_Event_Count | The number of attempts to transfer data from remapped sectors to spare sectors |
| | 197 | Current_Pending_Sector | Number of sectors waiting to be mapped |
| | 198 | Offline_Uncorrectable | The number of uncorrectable errors that occur when reading and writing sectors |
| | 199 | UDMA_CRC_Error_Count | CRC error count in Ultra DMA mode |
| | 200 | Multi_Zone_Error_Rate | The total number of errors occurring when writing data to sectors |

Table 3. SMART features.

“common risk” due to environmental factors. When analyzing hard disk failures, researchers can combine business scenario information with SMART data of hard disk to establish a multi-dimensional correlation model to improve the accuracy of prediction.

This article does not label the dataset, and researchers can set appropriate time windows according to their needs. The number of faulty and healthy disks in the file is shown in Table 5.

As can be seen from the table, the ratio of positive to negative samples for faulty disks and healthy disks is 146,784:712, indicating an imbalance in the sample. If it is not handled, a large number of hard disks will be judged as healthy disks by the established hard disk fault prediction model, so the faulty disks cannot be well detected. Researchers can improve the model by making improvements at both the data level and the algorithmic level, thereby achieving a better hard disk failure prediction model. The sample imbalance handling methods that can be adopted include, but are not limited to, the following:

- (1) Over-sampling methods: including SMOTE, Borderline-SMOTE, M-SMOTE, ADASYN and other methods, these methods select key small class samples according to different criteria to generate new samples.
- (2) Under-sampling methods: including Random under-sampling, under-sampling based on sample neighbor information, under-sampling based on clustering and other methods, which reduce the number of samples in large categories to achieve data balance.
- (3) Over-sampling and under-sampling combined methods: including SMOTE + ENN and SMOTE + Tomek, which take advantage of under-sampling to delete the noise data generated by over-sampling.
- (4) Cost-sensitive learning method: by adjusting the cost function of the model, the error classification of a few categories of samples has higher cost, and the deviation caused by category imbalance is limited.
- (5) Ensemble learning method: a certain number of samples are randomly selected from the categories with large sample size, and multiple models are trained by combining them with the categories with small sample size. Finally, voting or weighted voting method is used to generate classification prediction results.
- (6) Abnormal detection method: The focus of abnormal detection method is not to capture the differences between classes, but to find the characteristics of the class we are concerned with. Those that do not conform to these characteristics can be counted as another class.

Through the sample imbalance processing, we can obtain useful information in the case of huge imbalance between faulty disks and healthy disks, so as to achieve better hard disk fault prediction effect.

Technical Validation

Single collection experiment. To verify the rationality of the collection frequency used in this article. During testing, the open-source I/O stress testing tool FIO is used to perform fixed granularity (8KB) continuous writes on a single hard drive for 180 seconds, triggering the hard drive SMART acquisition every 8–12 seconds. Use the await metric in the iostat command as a basis for whether it causes interference. Await is the average wait time (in milliseconds) for each I/O operation, which includes the time spent in the kernel I/O queue and the time spent executing I/O on the storage device. Under the premise of fixed write granularity, there is almost no difference in the time for executing I/O on the hard disk. However, when collecting SMART data from the hard

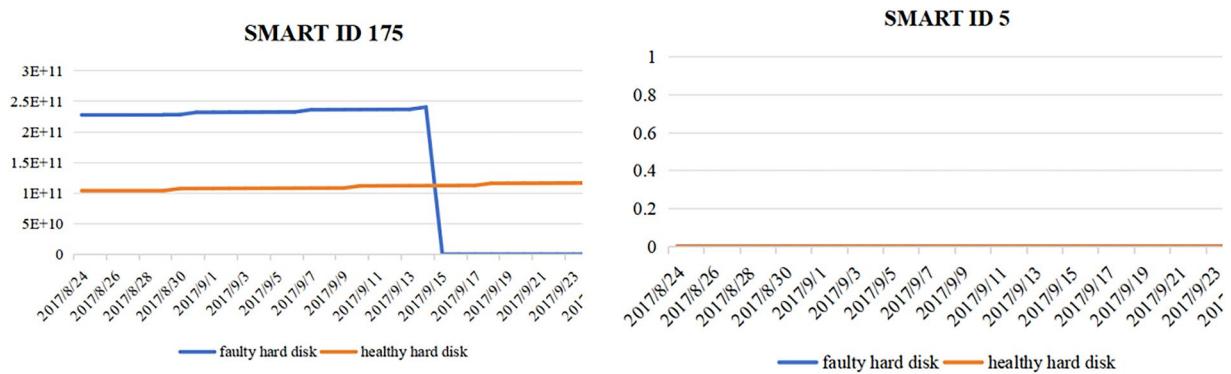


Fig. 2 Analysis of some SMART features. (a) Visualization of SMART ID 175. (b) Visualization of SMART ID 5.

| Bureau | | | |
|--------------------|------------------------|-------------------|---------------------|
| Anhui_DX_10012 | Hainan_LT_10093 | Jiangsu_YD_10151 | Shanghai_DX_10242 |
| Beijing_DX_10022 | Hebei_LT_10103 | Jiangxi_YD_10161 | Shanxi_LT_10253 |
| Beijing_LT_10023 | Hebei_YD_10101 | Liaoning_DX_10182 | Sichuan_DX_10262 |
| Belrus_DX_20002 | Henan_DX_10122 | Liaoning_LT_10183 | Sichuan_YD_10261 |
| Chongqing_DX_10032 | Henan_LT_10123 | Ningxia_DX_10202 | Thailand_True_10001 |
| Chongqing_LT_10033 | Hubei_DX_10132 | Ningxia_YD_10201 | Tianjin_LT_10273 |
| Fujian_DX_10042 | Hunan_DX_10142 | Qinghai_DX_10212 | Xinjiang_DX_10292 |
| Gansu_DX_10052 | Hunan_LT_10143 | Shaanxi_DX_10222 | Xinjiang_LT_10293 |
| Guangdong_YD_10061 | Hunan_YD_10141 | Shandong_LT_10233 | Xinjiang_YD_10291 |
| Guizhou_DX_10082 | Indonesia_Telkom_10001 | Shandong_YD_10231 | Zhejiang_DX_10322 |
| Hainan_DX_10092 | Jiangsu_DX_10152 | Shandong_YD_10234 | Zhejiang_YD_10321 |

Table 4. Bureau information.

| Total Hard Disks | Healthy Disks | Offline Disks |
|------------------|---------------|---------------|
| 147,496 | 146,784 | 712 |

Table 5. Hard disk ratio.

disk, the length of the I/O queue will increase. The test in this article is expected to cause a sudden increase in the await value at the moment of collection.

As shown in the Fig. 3, a single acquisition of the hard disk SMART did not cause read/write blocking or abnormalities, only causing momentary jitter in await, but the jitter level did not reach an unacceptable level. Therefore, in this article, a full disk scan is performed at 1-hour intervals every day, and data is collected at 10 second intervals. SMART information of 50–100 hard drives is collected each time, while avoiding busy periods of the front-end business system. This strategy is reasonable.

Collect different quantities experiment. As shown in Fig. 4, this article tested the performance jitter duration experienced by the storage system when collecting SMART data from hard drives. The traditional method of collecting SMART data from 10000 hard drives at once would cause a performance jitter of 6440 milliseconds, while the method proposed in this article can control this number within 100 milliseconds.

Interference experiment for video services. This experiment was conducted in ZTE's video business. The front-end business involved 20 clients playing 30 minutes of ultra clear video, and the normal state was that the storage system provided equal and constant data read bandwidth to all clients.

Figure 5(a) shows the case where no hard disk information is collected, Fig. 5(b) shows the case where all hard disks are collected at once using traditional methods, and Fig. 5(c) shows the case of the time-sharing hard disk information collection method in this article. The vertical axis of Fig. 5 represents the playback bandwidth of the client. From the test results, it can be seen that compared with Fig. 5(a), although the average bandwidth of Fig. 5(b) only decreased by 0.88% overall, the system was unable to serve continuously for 9.18 seconds during SMART collection. There was significant performance jitter in the first 11.21 seconds and the following 4.23 seconds, and users experienced severe lag for more than 20 seconds while watching live videos; Compared with Fig. 5(a), although the average bandwidth of Fig. 5(c) decreased by 0.96% overall, there was no system outage throughout the entire process, and users did not notice any abnormalities in video playback. From this, it can be seen that the work presented in this article has improved user experience and reduced interference with front-end business.

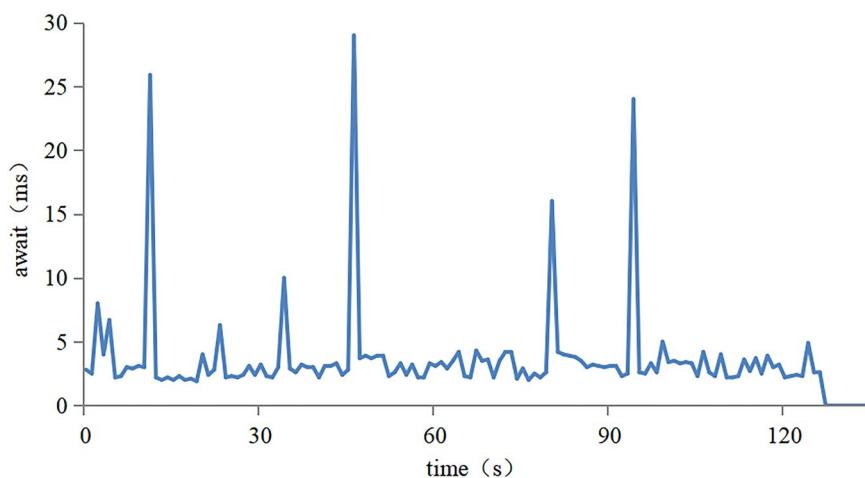


Fig. 3 Single collection of hard disk SMART causes slight jitter in await.

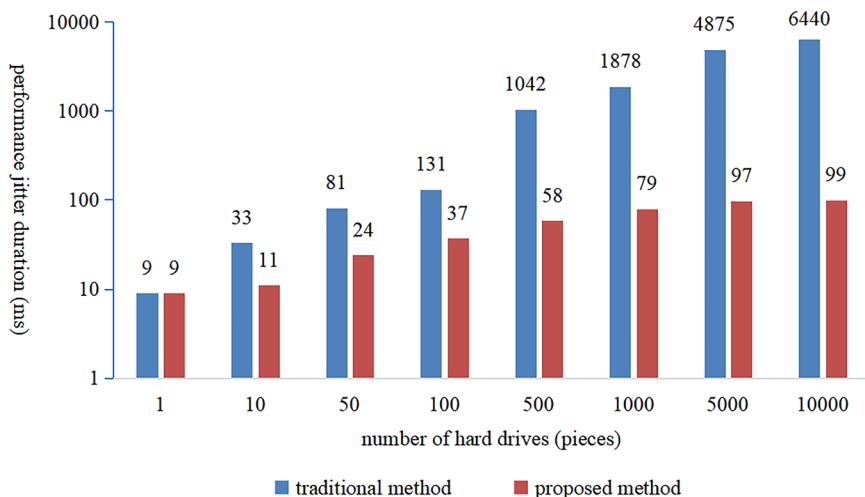


Fig. 4 Performance jitter test.

Data reliability. The dataset in this article is real-time collected results in a real environment. To verify the reliability of the dataset, we performed an integrity check on the `log_hdsmart_param.csv` file that stores SMART data in the dataset. For better comparison, we also conducted integrity checks on the SMART data of the ST4000DM000 hard drive model released by Backblaze in 2022. The results are shown in Tables 6, 7.

From the table, it can be seen that the blank values in this dataset are mainly concentrated in `thresh`, `type`, `update`, `value`, `worst`. However, the blank values only account for 10.7% of all data in the entire column and 5.3% of all data, indicating a relatively small number of blank data. And all the blank values mentioned above occur on `SAS_HDD` type hard drives. This is due to significant differences in underlying protocols, design objectives, and monitoring mechanisms between `SAS` disks and `SATA` hard drives, which result in different representations of SMART attributes. These values are not caused by data loss during collection or other reasons, but rather stem from the different properties of hard drive types. Therefore, compared to `SATA` disks, there will be blanks in `thresh`, `type`, `update`, `value`, and `worst`. We recommend discussing whether a hard drive is faulty separately for different types of hard drives, as their characteristic properties differ.

The Backblaze dataset has 911,132,360 blank values, accounting for 14.78% of all data, which is significantly higher than the SMART-Z dataset proposed in this paper. By comparing the data in Tables 6, 7, it can be more intuitively observed that the dataset collected in this article has better integrity, which can prove the feasibility of the collection technology.

Usage Notes

The SMART-Z dataset proposed in this article saves information in `.csv` file format and is publicly available on open source frameworks (<https://doi.org/10.17605/OSF.IO/24Y6G>). Convenient for researchers to download and process datasets. The SMART-Z dataset includes basic information of all hard drives, SMART attributes of hard drives, and basic information of all offline hard drives. Compared with other publicly available hard drive

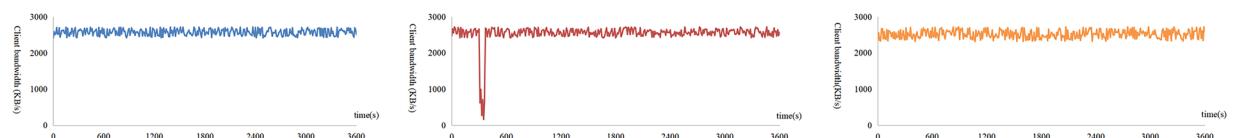


Fig. 5 Client bandwidth testing under different conditions for video services. **(a)** Do not collect hard disks SMART. **(b)** Method of collecting all hard disks SMART at once. **(c)** Proposed method.

| Column Name | Blank Count | Blank Ratio(%) |
|-------------|-------------|----------------|
| thresh | 20,953,799 | 10.7 |
| type | 20,953,799 | 10.7 |
| update | 20,953,799 | 10.7 |
| value | 20,953,799 | 10.7 |
| worst | 20,953,799 | 10.7 |
| keyHDid | 0 | 0 |
| keyScanTime | 0 | 0 |
| SMART_ID | 0 | 0 |
| SMART_NAME | 0 | 0 |
| rawvalue | 0 | 0 |
| Total | 104,768,995 | 5.3 |

Table 6. Integrity inspection.

| Column Name | Blank Count | Blank Ratio(%) |
|-------------------------|-------------|----------------|
| SMART Raw Or Normalized | 911,132,360 | 73.92 |
| model | 0 | 0 |
| serial_number | 0 | 0 |
| failure | 0 | 0 |
| date | 0 | 0 |
| Total | 911,132,360 | 14.78 |

Table 7. Backblaze integrity inspection.

datasets, this dataset adds multiple attributes such as critical values, worst-case values, device IP, business scenarios, and power on time of offline hard drives. In addition, through the observation in this article, it was found that when the SMART of the hard drive first showed abnormal performance, the median time to actual failure was 7 days, and the farthest was 56 days. We offer the following suggestions to researchers:

- (1) You can pay attention to the setting of time windows and the filtering of SMART features. Different models and manufacturers of faulty hard drives may experience abnormalities on different SMART.
- (2) In this article, we only declare all offline hard drives. There are many factors that can cause a hard drive to go offline, and a malfunction is just one of them. This needs to be discussed in detail in conjunction with the other two files, log_hdsmart_base.ase.csv and log_hdsmart_param.csv.
- (3) This article does not include a hard disk fault label, as label settings can have different impacts on the accuracy of hard disk fault prediction algorithms. Therefore, we only provide all offline disk information. You can set appropriate labels based on your own needs and experience. You can label the fault a few days in advance, or you may only set the fault for that day.
- (4) The business scenario information includes the physical deployment location information of the hard disk. Hard disks in the same business scenario may be affected by the same environmental factors, which can also be considered as factors for predicting hard disk failures. Similarly, IP addresses can also be considered.
- (5) The offline hdsmart.csv displays the cumulative power on duration of all offline disks. Long term power on may cause high temperature of the hard drive or continuous wear of the magnetic head. Therefore, it can also be explored whether the power on duration of the hard drive will have a certain impact on its lifespan.
- (6) This dataset annotates the warning levels (Old_age/Pre fail) of SMART attributes for different vendors. You can assign different weights to different SMART attribute warning levels, which provides a new direction for hard disk failure prediction.

Code availability

The code for collecting hard disk fault prediction datasets can be implemented in an open science framework (<https://doi.org/10.17605/OSF.IO/24Y6G>). Get it up. Shell scripts can be run in a Linux system environment.

Received: 22 April 2025; Accepted: 24 June 2025;
Published online: 02 July 2025

References

1. AllenBruce. Monitoring hard disks with smart. *Linux journal* 2004 (2004).
2. Friedman, J. H. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics* **29**, 1189–1232 (2001).
3. Li, J. *et al.* Hard drive failure prediction using Decision Trees. *Reliability Engineering & System Safety* **164**, 55–65 (2017).
4. Lu, S. *et al.* Making Disk Failure Predictions SMARTer. *File and Storage Technologies* 151–167 (2020).
5. Wei, S., Yang, H., Chen, Z. & Wang, P. SMART-Z (2017–2018) <https://doi.org/10.17605/OSF.IO/24Y6G> (2025).
6. Ma, A. *et al.* RAIDShield. *ACM Transactions on Storage* **11**, 1–28 (2015).

Acknowledgements

The authors would like to thank all the anonymous reviewers for their helpful comments and suggestions. This research was supported by a grant from the National Key R&D Program of China (2021YFB3101105), the Tianjin Manufacturing High Quality Development Special Foundation (20232185), and the ZTE Foundation(HC210412010101). We also thank Tianjin University of Technology for finding the innovation and entrepreneurship training program for university students.

Author contributions

Data Acquisition, ShutingWei, Hongzhang Yang, Zhengguang Chen and Ping Wang; Conceptualization, ShutingWei and Hongzhang Yang; Data curation, ShutingWei and Zhengguang Chen; Formal analysis, ShutingWei; Funding acquisition, Hongzhang Yang; Investigation, Ping Wang; Methodology, ShutingWei and Hongzhang Yang; Supervision, Hongzhang Yang; Roles/Writing - original draft, ShutingWei; and Writing - review & editing, Zhengguang Chen and Hongzhang Yang.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to H.Y.

Reprints and permissions information is available at 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