

VIDEO MEMORY MANAGEMENT METHOD, APPARATUS, DEVICE AND SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The application claims the benefits of priority to Chinese Application No. 202011219652.6, entitled "Video Memory Management Method, Apparatus, Device and System", filed on November 3, 2020, which is incorporated herein by reference in its entirety.

TECHNICAL FIELD

[0002] The present application relates to a technical field of machine learning, and in particular relates to a video memory management method, apparatus and system, a machine learning system, and an electronic device.

BACKGROUND

[0003] With the continuous development of deep learning algorithms and the support of the computing power of a graphics processing unit (GPU), deep learning has become a vital part of an enterprise product data flow. In order to support large-scale deep learning applications, an enterprise usually build a shared GPU cluster to support the development of products across a plurality of fields, such as computer vision, natural language processing, speech recognition, recommendation and advertising services, etc.

[0004] In order to improve a utilization rate of GPU resources and throughput of the entire GPU cluster, the deep learning system allows a plurality of deep learning tasks to be executed simultaneously on one GPU, so that more deep learning training tasks can be completed with the same amount of resources. At present, a typical method for reusing GPU memory resources is that video memory management is performed by a unified video memory allocator in a deep learning framework; when the allocator receives an application for video memory resources from any task, as long as a GPU executing the task has idle

video memory resources, a corresponding video memory space is allocated to the task, regardless of video memory resource requirements of other tasks that are simultaneously executed on the GPU. Such processing method may speed up a small-batch training speed of the tasks.

[0005] However, in the process of implementing the present invention, the inventors have found that the foregoing technical solution has at least the following problems: 1) The foregoing method for reusing resources does not provide any performance isolation guarantee, which will bring uncontrollable mutual influence between the plurality of tasks. Specifically, when the GPU is allocated to a "resource guarantee" task for separate use, a deep learning system may guarantee task training performance thereof. Due to the lack of a performance isolation mechanism of the GPU, if there are other tasks that are executed together on such GPU, potential competition for video memory resources may lead to serious performance degradation of the "resource guarantee" task. 2) With the proceeding of training, GPU video memory requirements of the "resource guarantee" task may suddenly increase, and if a GPU video memory is occupied by other tasks at this time, the "resource guarantee" task fails, which is even more unacceptable. To sum up, how to manage the shared video memory resources of the machine learning system, so as to improve the utilization rate of the GPU cluster under the premise of guaranteeing the usage of the video memory resources of the high-priority task, has become an urgent problem for those skilled in the art.

SUMMARY OF THE DISCLOSURE

[0006] The present application provides a video memory management method to solve the problem that the performance of a high-priority task cannot be guaranteed in the prior art. The present application further provides a video memory management apparatus and system, a machine learning system, and an electronic device.

[0007] The present application provides a video memory management method, including: determining priorities of a plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated for a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resources required by the high-priority task, releasing at least a part of video memory resources occupied by a low-priority task; and allocating the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[0008] Optionally, the method further includes: releasing idle video memory resources occupied by the plurality of machine learning tasks.

[0009] Optionally, the releasing idle video memory resources occupied by the plurality of machine learning tasks includes: determining the usage status information of video memory resources of the machine learning task; and if the information satisfies a video memory resource release condition, releasing the idle video memory resources.

[0010] Optionally, the usage status information includes: an upper limit of the video memory resources actually used by the task; and the release condition comprises: a duration during which the amount of the video memory resource allocated to the task is larger than the upper limit reaches a duration threshold.

[0011] Optionally, the method further includes: releasing idle video memory resources of other high-priority tasks.

[0012] Optionally, the releasing at least a part of the video memory resources occupied by a low-priority task includes: if the amount of idle video memory resources of the low-priority task is larger than or equal to the amount of the required video memory resources, releasing the idle video memory resources occupied by the low-priority task.

[0013] Optionally, the releasing at least a part of the video memory resources occupied by a low-priority task includes: if the amount of the idle video memory resources of the low-priority task is smaller than the amount of the required video memory resources, allocating internal memory resources to the low-priority task, and releasing at least a part of the video memory resources used by the low-priority task, so as to continue to execute the low-priority task at least according to tensor data in an internal memory space.

[0014] Optionally, the method further includes: if the amount of the allocatable video memory resources increases to the amount of the video memory resources required by the low-priority task, allocating the video memory resources to the low-priority task, so as to continue to execute the low-priority task according to the tensor data in the video memory space.

[0015] Optionally, the low-priority tasks include: iterative learning tasks; and the releasing the video memory resources used by the low-priority task includes: after the low-priority task completes current iterative learning, releasing at least a part of the used video memory resources.

[0016] Optionally, before the releasing at least a part of the video memory resources used by the low-priority task, the method further includes: allocating internal memory resources to the high-priority task, so as to execute the high-priority task according to the tensor data in the internal memory space and the tensor data in the video memory space.

[0017] Optionally, after the allocating video memory resources to the high-priority task, the method further includes: releasing the internal memory resources of the high-priority task.

[0018] Optionally, the machine learning tasks include distributed deep learning tasks.

[0019] The present application further provides a video memory management method, including: executing a machine learning task by a graphics processing unit; determining the usage status information of video memory resources of the machine learning task; and if the information satisfies a video memory resource release condition, releasing idle video memory resources occupied by the task, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[0020] The present application further provides a video memory management apparatus, including: a priority determination unit, configured to determine priorities of a plurality of machine learning tasks executed by a graphics processing unit; a video memory release unit, configured to release at least a part of video memory resources occupied by a low-priority task if video memory resources are to be allocated to a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resource required by the high-priority task; and a video memory allocation unit, configured to allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[0021] The present application further provides an electronic device, including: a processor and a memory, wherein the memory is configured to store a program for implementing a video memory management method, and after being powered on and executing the program of the method by the processor, the device performs the following steps: determining priorities of a plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated for a high-priority task, and an allocatable video memory space is smaller than a video memory space required by the high-priority task, releasing at least a part of a video memory space of a low-priority task;

and allocating the video memory space to the high-priority task according to the video memory space released from the low-priority task, so as to execute the high-priority task at least according to tensor data in a video memory.

[0022] The present application further provides a video memory management apparatus, including: a task executing unit configured to execute a machine learning task by a graphics processing unit; an information determining unit configured to determine the usage status information of video memory resources of the machine learning task; and a video memory release unit configured to release idle video memory resources occupied by the task if the information satisfies a video memory resource release condition, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[0023] The present application further provides an electronic device, including: a processor and a memory, wherein the memory is configured to store a program for implementing a video memory management method, and after being powered on and executing the program of the method by the processor, the device performs the following steps: executing a machine learning task by a graphics processing unit; determining the usage status information of video memory resources of the machine learning task; if the information satisfies a video memory resource release condition, releasing idle video memory resources occupied by the task, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[0024] The present application further provides a video memory management system, including: a storage resource coordinator, configured to determine priorities of a plurality of machine learning tasks that are executed by a graphics processing unit; and if video memory resources are to be allocated for a high-priority task, and the amount of

allocatable video memory resources is smaller than the amount of video memory resources of the high-priority task, send a video memory resource release instruction to a first storage resource allocator of low-priority tasks, and send a video memory resource allocation instruction to a second storage resource allocator of high-priority tasks; a first storage resource allocator configured to release at least a part of video memory resources occupied by the low-priority task according to the release instruction; a second storage resource allocator, configured to allocate the video memory resources to the high-priority task according to the allocation instruction, so as to execute the high-priority task at least according to tensor data in a video memory space.

[0025] Optionally, the allocator is further configured to send the usage status information of the video memory resources of the tasks; and the coordinator is further configured to send the video memory resource release instruction to the allocator if the information satisfies a video memory resource release condition.

[0026] Optionally, the allocator is further configured to send the information to the coordinator according to a preset period.

[0027] The present application further provides a machine learning system, including: a client configured to send the priority information of machine learning tasks to a server; a server configured to determine the priorities of the plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated to a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resources required by the high-priority task, release at least a part of video memory resources occupied by a low-priority task; and allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[0028] The present application further provides a computer-readable storage medium, wherein the computer-readable storage medium stores instructions, and when the instructions are executed on a computer, the computer is enabled to execute the foregoing various methods.

[0029] The present application provides a computer program product including instructions, wherein when the instructions are executed on a computer, the computer is enabled to execute the foregoing various methods.

[0030] Compared with the prior art, the present application has the following technical advantages:

[0031] According to the video memory management method provided by the embodiment of the present application, the priorities of the plurality of the machine learning tasks executed by the graphics processing unit are determined; if video memory resources are to be allocated for a high-priority task, and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, at least a part of the video memory resources occupied by the low-priority task are released; and the video memory resources are allocated to the high-priority task, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task when the allocatable video memory resources are insufficient, thereby realizing dynamic scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on the GPU, and in this way, the GPU video memory resources are allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

[0032] According to the video memory management method provided by the embodiment of the present application, the machine learning task is executed by the graphics processing unit; the usage status information of the video memory resources of the machine learning task is determined; if the information satisfies the video memory resources release condition, the idle video memory resources occupied by the task are released, so as to allocate the idle video memory resources to other machine learning tasks executed in parallel by the graphics processing unit. Due to such processing mode, the idle video memory resources occupied by the task are released in time. Therefore, a resource utilization rate of a whole cluster can be effectively enhanced.

[0033] According to the video memory management system provided by the embodiment of the present application, the priorities of the plurality of machine learning tasks executed by the graphics processing unit are determined through the storage resource coordinator; if the video memory resources are to be allocated to the high-priority task, and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, the video memory resource release instruction is sent to the first storage resource allocator of low-priority tasks, and the video memory resource allocation instruction is sent to the second storage resource allocator of high-priority tasks; the first storage resource allocator releases at least a part of the video memory resources occupied by the low-priority task according to the release instruction; the second storage resource allocator allocates the video memory resources to the high-priority task according to the allocation instruction, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task through the local storage resource coordinator of the GPU device when the allocatable video memory resources are insufficient, thereby realizing dynamic

scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on a GPU, and in this way, the GPU video memory resources can be allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

[0034] According to the machine learning system provided by the embodiment of the present application, the client sends the priority information of the machine learning tasks to the server; the server determines the priorities of the plurality of machine learning tasks executed by the graphics processing unit; if video memory resources are to be allocated to the high-priority task, and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, at least a part of the video memory resources occupied by the low-priority task are released; and the video memory resources are allocated to the high-priority task, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task when the allocatable video memory resources are insufficient, thereby realizing dynamic scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on the GPU, and in this way, the GPU video memory resources are allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0035] **FIG. 1** is a schematic flowchart of a video memory management method according to the present application.

[0036] **FIG. 2** is a schematic diagram showing an application scenario of a video memory management method according to the present application.

[0037] **FIG. 3** is a schematic diagram showing dynamic scaling of video memory resources of a video memory management method according to the present application.

[0038] **FIG. 4** is a schematic diagram showing a change in video memory resources of a video memory management method according to the present application.

[0039] **FIG. 5** is a schematic structural diagram of a video memory management apparatus according to the present application.

[0040] **FIG. 6** is a schematic structural diagram of a video memory management system according to the present application.

DETAILED DESCRIPTION

[0041] Many specific details are set forth in the description below to provide a thorough understanding of the present application. However, the present application can be implemented in many other ways that are different from those described here, and those skilled in the art may make similar promotions without violating the meaning of the present application. Therefore, the present application is not limited by the specific implementation disclosed below.

[0042] The present application provides a video memory management method, apparatus and system, a machine learning system, and an electronic device. Various solutions will be described in detail in combination with specific embodiments below.

[0043] **First Embodiment**

[0044] Please refer to **FIG. 1** which is a schematic flowchart of an embodiment of a video memory management method according to the present application. The method provided by this embodiment may include the following steps:

[0045] Step S101: determining priorities of a plurality of machine learning tasks executed by a graphics processing unit.

[0046] The video memory management method provided by the present application can be applied in a machine learning system, and is used for allocating and using GPU video memory resources when the plurality of machine learning tasks share the GPU video memory resources. The machine learning system may be a deep learning system constructed based on a deep learning computing framework, such as TensorFlow, PyTorch and other deep learning computing frameworks. The machine learning task, also called a machine learning model training task, may learn a machine learning model from training data. The model may be a model based on a deep neural network, and correspondingly, the machine learning task is a deep learning task. For example, the model is a named entity recognition model, a speech recognition model, a product recommendation model, or the like learned from the training data. The model may also be a machine learning model based on a non-neural network such as a decision tree.

[0047] To illustrate the video memory management method provided by the present application more intuitively, an application scenario of the method is described below first. As shown in **FIG. 2**, the machine learning system may be a distributed machine learning system, including one or more GPU-based computing nodes, also called GPU devices. Each GPU device may include one or more GPUs. A plurality of machine learning tasks may be executed in parallel on one GPU, and these machine learning tasks share video memory resources of the GPU. In addition, the GPU device further includes a central processing unit (CPU) and an internal memory, wherein the CPU may also be called a host

machine of the GPU. In **FIG. 2**, a node 1 includes a GPU1 and a GPU2, wherein a task D and a task E are simultaneously executed on the GPU1, and a task A, a task B and a task C are simultaneously executed on the GPU2.

[0048] According to the video memory management method provided by the present application, the video memory resources occupied by each task can be dynamically adjusted according to the performance guarantee priorities (referred to as priorities) of the plurality of tasks executed in parallel on the GPU. The higher the priority of the task, the more the performance of the task must be guaranteed. The priorities of the tasks may be determined according to application requirements, for example, only two priorities are set: a high priority and a low priority. For example, if a learning task 1 of the named entity recognition model is a "performance guarantee task" with a service level guarantee, and a learning task 2 of the speech recognition model is a "speculative execution task" without a service level guarantee, then the task 1 may be set with the high priority, and the task 2 may be set with the low priority.

[0049] During specific implementation, a plurality of priorities may also be set.

Table 1 shows priority setting information of machine learning tasks in an example.

Task ID	Task priority
Task 1 (Named Entity Recognition Model)	First level (highest priority)
Task 2 (Speech Recognition Model)	Second level (second highest priority)
Task 3 (product recommendation model)	Second level (second highest priority)
Task 4 (Language Model)	Third Level (lowest priority)
...	

Table 1. Machine Learning Task Table

[0050] It can be seen from Table 1 that three task priorities are set in this embodiment, wherein the first level is the highest priority, the second level is the second highest priority, and the third level is the lowest priority. In this case, if there is a first-level

task in parallel tasks, the performance of the first-level task must be guaranteed first. If the parallel tasks do not include a first-level task, the performance of a second-level task must be guaranteed.

[0051] Step S103: if video memory resources are to be allocated for a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resources required by the high-priority task, releasing at least a part of a video memory space occupied by a low-priority task.

[0052] In the video memory management method provided by the present application, when allocating the video memory resources to the high-priority task, if the amount of the allocatable video memory resources on the GPU is smaller than the amount of the video memory resources required by the high-priority task, a part or all of a video memory space occupied by the low-priority task should be released, and the released video memory space is allocated to the high-priority task, so as to store tensor data of the high-priority task in a video memory as much as possible, and thus the executing performance of the high-priority task can be guaranteed.

[0053] The amount of the required video memory resources may be the amount of video memory resources to be increased for a deep learning task being executed, or may be the amount of video memory resources initially required by a deep learning task to be executed. The amount of the required video memory resources may be determined by the deep learning task. Since a method for determining the amount of the required video memory resources belongs to a well-known prior art, it is not repeated here.

[0054] In this embodiment, if a certain deep learning task in the machine learning task table of one GPU is to be executed, the video memory resources required by the task is determined first. After determining the video memory resources required by the task, if the amount of the allocatable video memory resources of the GPU is smaller than the

amount of the video memory resource required by the task, the priority of the task and the priorities of other tasks that are executed on the GPU should be determined. If the priority of the task to be executed is higher than the priorities of the other tasks being executed, the video memory resources occupied by the low-priority task should be released to guarantee the performance of the high-priority task.

[0055] During specific implementation, the amount of the video memory resources that should be released may be determined according to the amount of the video memory resources required by the task to be executed. For example, if the high-priority task needs 1G video memory resources, 1.2 G video memory resources occupied by the low-priority task may be released.

[0056] During specific implementation, the video memory resources occupied by one or the plurality of low-priority tasks may be released. For example, if the high-priority task needs 5G video memory resources, and the video memory resources that are released from the low-priority task is insufficient, the video memory resources occupied by the plurality of low-priority tasks may be released to release the video memory resources more than 5G.

[0057] In an example, the method may further include the following steps: releasing idle video memory resources occupied by the plurality of machine learning tasks. During the implementation of the present invention, the inventors found that most deep learning tasks cannot fully utilize all of the allocated GPU memory resources at all times, and usually have idle video memory resources. Through research, the inventors have found that the idle video memory resources of the deep learning tasks may be caused by the following reasons.

[0058] 1) Product-oriented deep learning training tasks usually include a lot of computing parts, some of which are not easy to parallelize, so it is difficult to completely

occupy the GPU video memory resources, such as image sampling in image neural networks, feature extraction in advertisement models, data augmentation in computer vision, or the like.

[0059] 2) The deep learning system deals with massive training data. With the increase of data, a lot of time will be spent in a network data synchronization stage of the model in super-large distributed training. As shown in **FIG. 2**, the task E is executed on the node 1 and the node n-1 at the same time. Correspondingly, the GPU video memory resources are also idle during the data synchronization stage.

[0060] 3) Distributed deep learning training usually uses a synchronous stochastic gradient descent (SGD) method, which requires that the resources required by a training task need to be satisfied at the same time, and then the training task may start. Therefore, from the perspective of a cluster scheduler, when the resources are insufficient, the scheduler needs to reserve some available resources for the distributed tasks until all of the required resources are satisfied. Such reservation process also causes the GPU memory resources to be at an idle waiting state.

[0061] 4) In practical applications, some tensor data are only used in a specific deep learning training stage, such as data processing, model evaluation, or the like. These tensor data will not be used after the stage and will be cleared from the video memory, thereby producing idle video memory resources.

[0062] However, in an existing deep learning framework, the idle video memory resources will not be released, but will be always reserved for the task. The whole video memory resources occupied by the task are not reduced, and a part of the video memory resources are idle video memory resources. During subsequent executing processes, the task may use these video memory resources, e.g., for storing other tensor data, or may not

use these video memory resources, thereby results in a relatively low utilization rate of the GPU resources in many cases.

[0063] According to the method provided by this embodiment, based on the usage status of the video memory resources, the idle video memory resources that are not used by the tasks are released, and thus can be timely allocated to other tasks, to realize optimization of the processing of the plurality of tasks in a shared GPU resource scenario, and prevent the other tasks from queuing up. Therefore, a resource utilization rate of the GPU is improved, and then throughput of the shared GPU cluster can be improved.

[0064] During specific implementation, the releasing idle video memory resources occupied by the plurality of machine learning tasks may include the following sub-steps: 1) determining usage status information of video memory resources of the machine learning task; and 2) if the information satisfies a video memory resource release condition, releasing idle video memory resources. The usage status information includes: an upper limit of the video memory resources actually used by the task. The release condition includes: a duration during which the amount of video memory resources allocated to the task is larger than the upper limit reaches a duration threshold. For example, during the executing of the task, the upper limit value (peak value) of the video memory resources that are actually used is determined every 10 seconds. When the amount of the video memory resources occupied by the task is larger than the peak value that is actually needed for 30 consecutive seconds (after the usage status information has been determined three times), the idle video memory resources will be released.

[0065] In an example, the releasing at least a part of a video memory space occupied by a low-priority task may include the following sub-steps: if the amount of idle video memory resources of the low-priority task is larger than or equal to the amount of required video memory resources, releasing the idle video memory resources occupied by

the low-priority task. In this case, both types of tasks may be executed based on data stored in the video memory space, so that not only the executing performance of the high-priority task can be guaranteed, but also the performance of the low-priority task can be prevented from being affected.

[0066] In another example, the releasing at least a part of a video memory space occupied by a low-priority task may further include the following sub-steps: if the amount of idle video memory resources of the low-priority task is smaller than the amount of required video memory resources, allocating internal memory resources to the low-priority task, and releasing at least a part of video memory resources used by the low-priority task, so as to continue to execute the low-priority task at least according to tensor data in an internal memory space.

[0067] The video memory resources used by the low-priority task are not idle video memory resources, and the tensor data of the low-priority tasks is still stored in these video memory resources which are needed by the low-priority task needs to guarantee the executing performance. However, a part or all of the video memory resources used by the low-priority task are released, and a part or all of the tensor data of the low-priority task is temporarily transferred from the video memory space to the internal memory space for storage, so that the low-priority task can be executed continuously at least according to the tensor data in the internal memory space. If only a part of the video memory resources used by the low-priority task are released, the low-priority task can be executed continuously based on one part of the tensor data in the internal memory space and another part of the tensor data in the video memory space. If all of the video memory resources used by the low-priority task are released, the low-priority task can be executed continuously based on all of the tensor data in the internal memory space. Hence, no matter what the situation is, the task will not fail.

[0068] By using such processing method, if video memory resource requirements of the high-priority task still cannot be met after the idle video memory resources of the low-priority task are released, the video memory resources occupied by the tensor data of the low-priority task need to be further released, and all or part of tensor data of the low-priority task are transferred to the internal memory space of the host machine of the GPU for storage, to execute the low-priority task at least based on the tensor data in the internal memory space. Therefore, the low-priority task and the high-priority task can be prevented from competing for the video memory resources, and the performance of the high-priority task can be guaranteed, at the cost of sacrificing part of the performance of the low-priority task.

[0069] As shown in **FIG. 3**, a sub-figure (a) shows video memory resources occupied by a deep learning task in an initial stage, including a tensor of data, wherein the height of a dotted line indicates a level of the occupied video memory resources. A sub-figure (b) shows that an upper limit of video memory usage increases with a corresponding increase of tensor data required by the task. These tensor data may be cached in a pool of an internal memory allocator (a video memory resource allocator) of the task, that is, cached in a video memory space allocated to the task, so that next small-batch tensor data of the task can be reused. The enlarged video memory resource of the task may accommodate three tensors of data. With continuous executing of the task, some tensor data are only used in a specific deep learning training stage, and these tensor data will not be used after the stage and will be cleared from the video memory, thereby causing idle video memory resources. A sub-figure (c) shows that after one tensor of data is cleared, idle video memory resources appear, while the level of the video memory resources does not change.

[0070] According to the method provided by the embodiment of the present application, the upper limit of video memory usage is dynamically adjusted according to the video memory resources actually required by the task. In this embodiment, the video memory resources currently in use are actively detected, and the idle video memory resources are released, thereby adjusting the upper limit of video memory usage to an appropriate value. A sub-figure (d) shows the level of the video memory resources after the available video memory (idle video memory resources) is recycled (released), after the foregoing available video memory has not been used for a period of time.

[0071] In this embodiment, a deep learning system open an interface to the outside, which allows the upper limit of GPU video memory usage for the task to be raised or reduced when the task is executed, and even reduced to a value lower than an actual requirement of the task. A sub-picture (e) may show a situation that when there is a high-priority task and the allocatable video memory is insufficient, a tensor of data of a low-priority task is transferred to an internal memory space of a host machine, so that the released video memory resources can be allocated to the high-priority task, to guarantee the performance of the high-priority task. By using such processing method, even if the video memory resources are insufficient, the low-priority task can still be executed continuously without failing.

[0072] In an example, the method may further include the following steps: if the amount of the allocatable video memory resources increases to the amount of the video memory resources required by the low-priority task, allocating the video memory resources to the low-priority task, so as to continue to execute the low-priority task according to the tensor data in the video memory space. By using such processing method, when the GPU video memory resources are no longer in shortage, the upper limit of the video memory for the low-priority task may be raised, and the tensor of the low-priority task may also be

reallocated on the GPU. A sub-figure (f) shows a situation that when the GPU video memory resources are no longer in shortage, the tensor data of the low-priority task is transferred from the internal memory space to the video memory space again, so as to improve the performance of the low-priority task as much as possible.

[0073] In an example, the low-priority tasks include: iterative learning tasks. The releasing the video memory resources used by the low-priority task may be implemented in the following manner: after the low-priority task completes current iterative learning, releasing the used video memory resources.

[0074] In a deep learning scenario, to perform a gradient descent algorithm on the training set, an entire data set is usually divided into several small training sets, and one small subset is trained each time. In this way, on the one hand, the problem of a huge amount of calculation caused by all of the data sets participating in training at one time can be avoided; and on the other hand, there is no big difference between a gradient direction of the subset and a gradient of the whole data set, which can guarantee the correctness of training. Such training method is also called a small batch training method. Such learning task is called the iterative learning task. Every time a small-batch training is performed, various data involved in the training must be stored in a storage space, and a container for storing the data is called a tensor (Tensor). The tensor is a data unit for storing data in a deep learning framework. As a data container, the tensor may store data during training, and a task may include a plurality of tensors in one training. Tensor data may be an arbitrary quantity of multi-dimensional arrays composed of a set of original values, including training data for one small-batch training, model parameters generated during training, and data of various intermediate nodes in a network, or the like.

[0075] In this embodiment, by transferring the tensor data between the video memory and the internal memory, the utilization rate of the GPU video memory resources

is improved under the premise of guaranteeing the executing performance of the high-priority task. By using the processing method of transferring the tensor data between the video memory and the internal memory, the internal memory resources can be used as the video memory resources when the video memory resources are insufficient, but this usually causes huge data copy overheads.

[0076] By studying the unique characteristics of the deep learning tasks, that is, the characteristics of the foregoing iterative learning, the inventors of the present invention have found that in the deep learning tasks, training is performed in small batches, and the tensor data will be created and destroyed in small batches. The same tensor will be repeatedly created between small batches, so as to choose to dynamically adjust the upper limit of video memory usage for the task on the boundary of the small batches. At this time, the tensor data has been released, so that data copy explicitly performed between the video memory and the internal memory can be avoided, thereby avoiding huge copy overheads.

[0077] During specific implementation, before the releasing the video memory resources used by the low-priority task, the method may further include the following step: allocating a part of internal memory resources and a part of video memory resources to the high-priority task, so as to execute the high-priority task according to the tensor data in the internal memory space and the tensor data in the video memory space.

[0078] During specific implementation, after allocating video memory resources to the high-priority task, the method may further include the following step: releasing the internal memory resources of the high-priority task, so that the high-priority task can be executed completely based on the tensor data in the video memory space.

[0079] As shown in **FIG. 4(a)**, video memory resources required by a high-priority task suddenly increase at time T₀, the required video memory resources are denoted as a.

At this time, allocatable video memory resources are insufficient, so that a is first occupied in an internal memory space, and the high-priority task is executed temporarily based on the tensor data in an internal memory. At time T1, after a small-batch training of a low-priority task is completed, tensor data of the low-priority task is transferred to the internal memory, and video memory resources a released from the low-priority task are allocated to the high-priority task, and meanwhile internal memory resources a occupied by the high-priority task are released, so that the video memory resources of the high-priority task are increased by a, and the high-priority task is continuously executed on the video memory. By using such processing method, the performance of the high-priority task is only reduced in an extremely short period of time, and cannot be failed. Meanwhile, huge copy overheads may be avoided.

[0080] As shown in **FIG. 4(b)**, during the time period from T0 to T1, video memory resources b used by a low-priority task is released, and b is occupied in an internal memory space. The low-priority task is executed based on tensor data in an internal memory space, and thus part of the operating performance is lost. Then, after video memory resources are released from other tasks, the video memory resources of the low-priority task at time T1 are increased by b, and the low-priority task is continuously executed based on tensor data in a video memory space.

[0081] In one example, not only a video memory space occupied by the low-priority task can be released, but also an idle video memory space occupied by other high-priority tasks can be released. By using such processing method, more video memory resources may be released, and the performance of the low-priority task are effectively improved.

[0082] In an example, the method may further include S105: according to the video memory space released from the low-priority task, allocating the video memory space to

the high-priority level task so as to execute the high-priority task at least according to tensor data in a video memory.

[0083] It can be seen from the foregoing embodiment, according to the video memory management method provided by the embodiments of the present application, the priorities of the plurality of machine learning tasks executed by the graphics processing unit are determined; if the video memory resources are to be allocated to the high-priority task, and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, at least a part of the video memory resources occupied by the low-priority task are released; and the video memory resources are allocated to the high-priority task, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task when the allocatable video memory resources are insufficient, thereby realizing dynamic scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on the GPU, and in this way, the GPU video memory resources can be allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

[0084] **Second Embodiment**

[0085] In the foregoing embodiments, a video memory management method is provided, and correspondingly, the present application further provides a video memory management apparatus. The apparatus corresponds to the foregoing method embodiments. Since the apparatus embodiments are basically similar to the method embodiments, the apparatus embodiments are simply described, and for relevant parts, reference can be made

to part of the description of the method embodiments. The apparatus embodiments described below are illustrative only.

[0086] Please refer to **FIG. 5** which is a schematic structural diagram of an embodiment of a video memory management apparatus according to the present application. The present application further provides a video memory management apparatus, including:

[0087] a priority determination unit 501, configured to determine priorities of a plurality of machine learning tasks executed by a graphics processing unit;

[0088] a video memory release unit 502, configured to release at least a part of video memory resources occupied by a low-priority task if video memory resources are to be allocated for a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resource required by the high-priority task; and

[0089] a video memory allocation unit 503, configured to allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[0090] **Third Embodiment**

[0091] In the foregoing embodiments, a video memory management method is provided, and correspondingly, the present application further provides an electronic device. The device corresponds to the foregoing method embodiments. Since the device embodiment is basically similar to the method embodiments, the device embodiments are simply described, and for relevant parts, reference can be made to part of the description of the method embodiments. The device embodiments described below are illustrative only.

[0092] The electronic device provided by this embodiment includes a processor and a memory, wherein the memory is configured to store a program for implementing a video

memory management method, and after being powered on and executing the program of the method by the processor, the device performs the following steps: determining priorities of a plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated for a high-priority task, and an allocatable video memory space is smaller than a video memory space required by the high-priority task, releasing at least a part of a video memory space of a low-priority task; and allocating the video memory space to the high-priority task according to the video memory space released from the low-priority task, so as to execute the high-priority task at least according to tensor data in a video memory.

[0093] **Fourth Embodiment**

[0094] Corresponding to the foregoing video memory management method, the present application further provides a video memory management method. The parts in this embodiment that are the same as those in the first embodiment will not be described again, please refer to the corresponding parts in the first embodiment.

[0095] In this embodiment, the method may include the following steps:

[0096] Step 1: Executing a machine learning task by a graphics processing unit.

[0097] Step 2: Determining the usage status information of video memory resources of the machine learning task.

[0098] Step 3: If the information satisfies a video memory resource release condition, releasing idle video memory resources occupied by the task, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[0099] The machine learning tasks include, but are not limited to, deep learning tasks. The idle video memory resources may include at least one of the following resources: idle video memory resources generated by modules that cannot be processed in

parallel in the deep learning task, and idle video memory resources generated when a plurality of resources required by the deep learning task are satisfied.

[00100] The machine learning task may be a distributed deep learning task. The idle video memory resources may be idle video memory resources generated when a plurality of image processing units corresponding to a distributed deep learning task synchronize data.

[00101] Specifically, the idle video memory resources of the deep learning task may be caused by the following reasons.

[00102] 1) Idle video memory resources generated by the modules that cannot be processed in parallel in the deep learning task. Product-oriented deep learning training tasks usually include a lot of computing parts, some of which are not easy to parallelize, so it is difficult to completely occupy GPU memory resources, such as image sampling in image neural networks, feature extraction in advertisement models, data augmentation in computer vision, or the like.

[00103] 2) Idle video memory resources generated when the plurality of image processing units corresponding to the distributed deep learning task synchronize data. The deep learning system deals with massive training data. With the increase of data, a lot of time will be spent in a network data synchronization stage of the model in super-large distributed training. As shown in **FIG. 2**, a task E is executed in a node 1 and a node n-1 at the same time, and correspondingly, GPU memory resources are also idle during the data synchronization period.

[00104] 3) Idle video memory resources generated while waiting for the plurality of resources required by the deep learning task to be satisfied. Distributed deep learning training usually uses a synchronous stochastic gradient descent (SGD) method, which requires that the resources required by a training task need to be satisfied at the same time,

and then the training task may start. Therefore, from the perspective of a cluster scheduler, when the resources are insufficient, the scheduler needs to reserve some available resources for the distributed tasks until all of the required resources are satisfied. Such reservation process also causes the GPU memory resources to be at an idle waiting state.

[00105] 4) In practical applications, some tensor data are only used in a specific deep learning training stage, such as data processing, model evaluation, or the like, and these tensor data will not be used after the stage and will be cleared from the video memory, thereby producing idle video memory resources.

[00106] The usage status information includes, but is not limited to: an upper limit of the video memory resources actually used by the task. The release condition includes, but is not limited to: a duration during which the amount of video memory resources allocated to the task is larger than the upper limit reaches a duration threshold. The duration threshold may be determined according to application requirements, for example, set to 30 seconds.

[00107] For example, during the executing of the task, the upper limit value (peak value) of the video memory resources that are actually used is determined every 10 seconds. When the amount of the video memory resources occupied by the task is larger than the peak value that is actually needed for 30 consecutive seconds (after the usage status information is determined three times), the idle video memory resources will be released.

[00108] It can be seen from the foregoing embodiments, according to the video memory management method provided by the embodiment of the present application, the machine learning task is executed by the graphics processing unit; the usage status information of the video memory resources of the machine learning task is determined; if the information satisfies a video memory resources release condition, the idle video

memory resources occupied by the task are released, so as to allocate the idle video memory resources to other machine learning tasks executed in parallel by the graphics processing unit. Due to such processing method, the idle video memory resources occupied by the task are released in time. Therefore, a resource utilization rate of a whole cluster can be effectively improved.

[00109] **Fifth Embodiment**

[00110] In the foregoing embodiments, a video memory management method is provided, and correspondingly, the present application further provides a video memory management apparatus. The apparatus corresponds to the foregoing method embodiment. Since the apparatus embodiment is basically similar to the method embodiment, the apparatus embodiment is simply described, and for relevant parts, reference can be made to part of the description of the method embodiment. The apparatus embodiment described below is illustrative only.

[00111] In this embodiment, the apparatus includes:

[00112] a task executing unit, configured to execute a machine learning task by a graphics processing unit;

[00113] an information determining unit, configured to determine the usage status information of video memory resources of the machine learning task; and

[00114] a video memory release unit, configured to release idle video memory resources occupied by the task if the information satisfies a video memory resource release condition, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[00115] **Sixth Embodiment**

[00116] In the foregoing embodiment, a video memory management method is provided, and correspondingly, the present application further provides an electronic

device. The device corresponds to the foregoing method embodiment. Since the device embodiment is basically similar to the method embodiment, the device embodiment is simply described, and for relevant parts, reference can be made to part of the description of the method embodiment. The device embodiment described below is illustrative only.

[00117] The electronic device provided by this embodiment includes a processor and a memory, wherein the memory is configured to store a program for implementing a video memory management method, and after being powered on and executing the program of the method by the processor, the device performs the following steps: executing a machine learning task by a graphics processing unit; determining the usage status information of video memory resources of the machine learning task; if the information satisfies a video memory resource release condition, releasing idle video memory resources occupied by the task, so as to allocate the idle video memory resources to other machine learning tasks executed in parallel by the graphics processing unit.

[00118] **Seventh Embodiment**

[00119] Corresponding to the foregoing video memory management method, the present application further provides a video memory management system. The parts in this embodiment that are the same as those in the first embodiment will not be described again, please refer to the corresponding parts in the first embodiment.

[00120] Please refer to **FIG. 6** which is a schematic structural diagram of an embodiment of a video memory management apparatus according to the present application. The video memory management system provided by the present invention includes: a storage resource coordinator, a storage resource allocator (a first storage resource allocator) of low-priority tasks (such as a task A), and a storage resource allocator (a second storage resource allocator) of low-priority tasks (such as a task B).

[00121] The system provided by this embodiment realizes adaptive dynamic scaling optimization of GPU video memory resources for machine learning tasks through co-design of the storage resource coordinator and a machine learning computing framework. As shown in **FIG. 6**, the storage resource coordinator supports adaptive GPU video memory resource adjustment, and can be deployed in a GPU computing node to schedule and manage memory resources of one or more GPUs in the node, to dynamically adjust the usage of the GPU video memory resources for a plurality of tasks on one GPU. The machine learning computing framework corresponds to machine learning tasks, and each task is executed by the machine learning computing framework.

[00122] The machine learning computing framework may be an end-to-end machine learning platform, and various machine learning frameworks can have own ecosystems, which include various tools, libraries and other resources, and can help developers easily build and deploy applications supported by machine learning. In this embodiment, the machine learning computing framework is a deep learning computing framework, including but not limited to TensorFlow, PyTorch, MXNet, Caffe, etc.

[00123] The storage resource coordinator is configured to determine priorities of a plurality of machine learning tasks executed by a graphics processing unit. If video memory resources are to be allocated to a high-priority task, and the amount of allocatable video memory resources is smaller than the amount of video memory resources required by the high-priority task, a video memory resource release instruction is sent to the storage resource allocator of low-priority tasks, and a video memory resource allocation instruction is sent to the storage resource allocator of low-priority tasks (a second storage resource allocator). The storage resource allocator of low-priority tasks is configured to release at least a part of video memory resources occupied by a low-priority task according to the release instruction. The storage resource allocator of high-priority tasks is configured to

allocate the video memory resources to the high-priority task according to the allocation instruction, so as to execute the high-priority task at least according to tensor data of a video memory space.

[00124] In an example, the allocator is further configured to send the usage status information of the video memory resources of the machine learning tasks. During specific implementation, the allocator can send the information to the coordinator according to a preset period. Correspondingly, the coordinator is further configured to send a video memory resource release instruction to the allocator if the information satisfies a video memory resource release condition.

[00125] It can be seen from the foregoing embodiment, according to the video memory management system provided by the embodiments of the present application, the priorities of the plurality of machine learning tasks executed by the graphics processing unit are determined by the storage resource coordinator. If video memory resources are to be allocated for the high-priority task and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, the video memory resource release instruction is sent to the first storage resource allocator of low-priority tasks, and the video memory resource allocation instruction is sent to the second storage resource allocator of low-priority tasks; the first storage resource allocator releases at least a part of the video memory resources occupied by the low-priority task according to the release instruction; the second storage resource allocator allocates the video memory resources to the high-priority task according to the allocation instruction, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task by the local storage resource coordinator of the GPU device when the allocatable video memory resources are

insufficient, thereby realizing dynamic scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on the GPU, and in this way, the GPU video memory resources can be allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

[00126] Eighth Embodiment

[00127] Corresponding to the foregoing video memory management method, the present application further provides a machine learning system. The parts in this embodiment that are the same as those in the first embodiment will not be described again, please refer to the corresponding parts in the first embodiment. The machine learning system provided by the present application includes a client and a server.

[00128] The client is configured to send the priority information of machine learning tasks to the server. The server is configured to determine the priorities of the plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated for a high-priority task and the amount of allocatable video memory resources is smaller than the amount of video memory resources required by the high-priority task, release at least a part of the video memory resources occupied by a low-priority task; and allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[00129] The client includes, but is not limited to, mobile communication devices, that is, mobile phones or smart phones, as well as personal computers, personal digital assistants (PADs), iPads and other terminal device. The server may execute the machine learning tasks on a GPU cluster.

[00130] In this embodiment, a task service apparatus may be provided for a user via the client, and the user may determine a machine learning task to be executed by the task service apparatus of the client, and set the priority of the task, for example, selecting "second level" priority, and different service charges may need to be paid for different priorities. After determining the machine learning task to be executed and setting the priority of the task, the client may submit a task executing request to the server. In response to the request, the server may store the priority information of the task and store the task in a task table.

[00131] In this embodiment, the server may include one or more GPU computing nodes, each node may execute the machine learning task in a machine learning framework, and a storage resource coordinator deployed in the GPU computing node may obtain the priority information of the plurality of tasks that are to be executed by the computing node.

Table 2 shows machine learning task information in this embodiment.

Task ID	User ID	Task priority
Task 1 (Named Entity Recognition Model)	User A	First level (Performance Guarantee Tasks)
Task 2 (Named Entity Recognition Model)	User B	Second level (speculative execution tasks)
Task 3 (product recommendation model)	User C	Second level (speculative execution tasks)
Task 4 (Language Model)	User C	First level (Performance Guarantee Tasks)
...		

Table 2. Machine Learning Task Table

[00132] It can be seen from **Table 2** that the learning task 1 of the named entity recognition model of the user A is a "performance guarantee task", the learning task 2 of the named entity recognition model of the user B is a "speculative execution task", and the

priority of the "performance guarantee task" is higher than the priority of the "speculative execution task".

[00133] After the priorities of the plurality of machine learning tasks executed by the graphics processing unit are determined, the following operations can be performed: if video memory resources are to be allocated to the high-priority task, and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, releasing at least a part of the video memory resources occupied by the low-priority task; and allocating the video memory resources to the high-priority task, so as to execute the high-priority task at least according to the tensor data in the video memory space.

[00134] In an example, the server is further configured to determine the performance information of the machine learning tasks; and adjust the priorities of the tasks according to the performance information.

[00135] For example, the priority of the task A is originally the second level, but the performance of the task does not meet service level requirements of the "second level" required by the user by the system, so that the priority of the task may be adjusted to the first level, and thus the actual performance of the task can meet the service level requirements of the "second level" required by the user.

[00136] During specific implementation, the server can record the change information of the video memory resources during the executing of the machine learning task, and adjust the priority information of the task according to the change information. For example, if a high-priority task is executed 30% of the time based on tensor data in an internal memory space, the performance information of the task does not meet the service level requirements, and therefore, the priority of the task can be increased.

[00137] In an example, the server is further configured to determine the performance information of the machine learning tasks; and determine the priorities of the machine learning tasks that can be served.

[00138] For example, if the service level requirements of the task A and the task B can be met by the system, but the service level requirement of the task C cannot be met, a video memory resource management service can be provided for the task A and the task B.

[00139] It can be seen from the foregoing embodiments, according to the machine learning system provided by the embodiments of the present application, the client sends the priority information of the machine learning tasks to the server; the server determines the priorities of the plurality of machine learning tasks executed by the graphics processing unit; if video memory resources are to be allocated for the high-priority task and the amount of the allocatable video memory resources is smaller than the amount of the video memory resources required by the high-priority task, at least a part of the video memory resources occupied by the low-priority task are released; and the video memory resources are allocated to the high-priority task, so as to execute the high-priority task at least according to the tensor data in the video memory space. Such processing mode allocates the video memory resources occupied by the low-priority task to the high-priority task when the allocatable video memory resources are insufficient, thereby realizing dynamic scaling optimization of the GPU video memory resources occupied by the plurality of machine learning tasks executed in parallel on the GPU, and in this way, the GPU video memory resources are allocated to the other tasks under the premise of guaranteeing the performance of the high-priority task. Therefore, a resource utilization rate of a whole cluster can be improved, and meanwhile the performance of the high-priority task is guaranteed.

[00140] Although the present application is disclosed as the preferred embodiments above, it is not intended to limit the present application. Any person skilled in the art may make possible changes and modifications without departing from the spirit and scope of the present application. Therefore, the scope of protection of the present application should be defined by the claims of the present application.

[00141] In a typical configuration, a computing device includes one or more processors (CPUs), an input/output interface, a network interface, and a memory.

[00142] The memory may include a volatile memory, a random access memory (RAM), and/or a non-volatile memory in computer-readable media, e.g., a read-only memory (ROM) or a flash RAM. The memory is an example of the computer-readable medium.

[00143] 1. The computer-readable medium includes permanent and non-permanent, removable and non-removable media, and can store information by any method or technology. The information may be computer-readable instructions, data structures, program modules, or other data. Examples of computer storage media include, but are not limited to, a phase change memory (PRAM), a static random access memory (SRAM), a dynamic random access memory (DRAM), other type of random access memory (RAM), a read-only memory (ROM), an electrically erasable programmable read-only memory (EEPROM), a flash memory or other memory technologies, a compact disc read-only memory (CD-ROM), a digital versatile disc (DVD) or other optical storage, a cassette magnetic tape, tape and disk storage or other magnetic storage devices, or any other non-transmission media, which may be configured to store information accessible by a computing device. As defined herein, the computer-readable medium does not include transitory computer-readable media (transitory media), such as a modulated data signal and a carrier.

[00144] 2. Those skilled in the art should understand that the embodiments of the present application can be provided as a method, a system, or a computer program product. Therefore, the present application may be in the form of a pure hardware embodiment, a pure software embodiment, or an embodiment combining software and hardware. Moreover, the present application can be in the form of a computer program product implemented on one or more computer-usable storage media (including, but not limited to, a magnetic disk memory, a CD-ROM, an optical memory, and the like) including computer-usable program code.

WHAT IS CLAIMED IS:

[Claim 1] A video memory management method, comprising:
determining priorities of a plurality of machine learning tasks executed by a graphics processing unit;
if video memory resources are to be allocated for a high-priority task, and an amount of allocatable video memory resources is smaller than an amount of video memory resources required by the high-priority task, releasing at least a part of video memory resources occupied by a low-priority task; and
allocating video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[Claim 2] The method according to claim 1, further comprising:
releasing idle video memory resources occupied by the plurality of machine learning tasks.

[Claim 3] The method according to claim 2, wherein releasing the idle video memory resources occupied by the plurality of machine learning tasks comprises:
determining usage status information of video memory resources of the plurality of machine learning tasks; and
if the information satisfies a video memory resource release condition, releasing the idle video memory resources.

[Claim 4] The method according to claim 3, wherein
the usage status information comprises: an upper limit of video memory resources actually used by the task; and
the release condition comprises: a duration during which an amount of video memory resource allocated to the task is larger than the upper limit reaches a duration threshold.

[Claim 5] The method according to claim 1, further comprising:
releasing idle video memory resources of other high-priority tasks.

[Claim 6] The method according to claim 1, wherein releasing at least a part of the video memory resources occupied by the low-priority task comprises:

if an amount of idle video memory resources of the low-priority task is larger than or equal to an amount of the required video memory resources, releasing the idle video memory resources occupied by the low-priority task.

[Claim 7] The method according to claim 1 or 6, wherein releasing at least a part of the video memory resources occupied by the low-priority task comprises:

if an amount of the idle video memory resources of the low-priority task is smaller than an amount of the required video memory resources,

allocating internal memory resources to the low-priority task; and

releasing at least a part of the video memory resources used by the low-priority task, so as to continue to execute the low-priority task at least according to tensor data in an internal memory space.

[Claim 8] The method according to claim 7, further comprising:

if the amount of the allocatable video memory resources increases to an amount of the video memory resources required by the low-priority task, allocating video memory resources to the low-priority task, so as to continue to execute the low-priority task according to the tensor data in the video memory space.

[Claim 9] The method according to claim 7, wherein the low-priority task comprises an iterative learning task; and releasing the video memory resources used by the low-priority task comprises:

after the low-priority task completes current iterative learning, releasing at least a part of the video memory resources used by the low-priority task.

[Claim 10] The method according to claim 9, wherein before the releasing at least a part of the video memory resources used by the low-priority task, the method further comprises:

allocating internal memory resources to the high-priority task, so as to execute the high-priority task according to tensor data in the internal memory space and the tensor data in the video memory space.

[Claim 11] The method according to claim 10, wherein after allocating memory resources to the high-priority task, the method further comprises:

releasing the internal memory resources of the high-priority task.

[Claim 12] The method according to claim 1, wherein the machine learning tasks comprise distributed deep learning tasks.

[Claim 13] A video memory management method, comprising:

running a machine learning task by a graphics processing unit;

determining usage status information of video memory resources of the machine learning task; and

if the information satisfies a video memory resource release condition, releasing idle video memory resources occupied by the task, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[Claim 14] A video memory management apparatus, comprising:

a priority determination unit configured to determine priorities of a plurality of machine learning tasks executed by a graphics processing unit;

a video memory release unit configured to release at least a part of video memory resources occupied by a low-priority task if video memory resources are to be allocated for

a high-priority task, and an amount of allocatable video memory resources is smaller than an amount of video memory resource required by the high-priority task; and

a video memory allocation unit configured to allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.

[Claim 15] An electronic device, comprising a processor and a memory, wherein the memory is configured to store a program for implementing the video memory management method according to any one of claims 1 to 12, and the device is powered on and executes the program of the method by the processor.

[Claim 16] A video memory management apparatus, comprising:

a task executing unit configured to execute a machine learning task by a graphics processing unit;

an information determining unit configured to determine usage status information of video memory resources of the machine learning task; and

a video memory release unit configured to release idle video memory resources occupied by the task if the information satisfies a video memory resource release condition, so as to allocate the idle video memory resources to other machine learning tasks that are executed in parallel by the graphics processing unit.

[Claim 17] An electronic device, comprising:

a processor and a memory, wherein

the memory is configured to store a program for implementing the video memory management method according to claim 13, and the device is powered on and executes the program of the method by the processor.

[Claim 18] A video memory management system, comprising:

a storage resource coordinator configured to:

determine priorities of a plurality of machine learning tasks that are executed by a graphics processing unit; and

if video memory resources are to be allocated for a high-priority task, and an amount of allocatable video memory resources is smaller than an amount of video memory resources of the high-priority task,

send a video memory resource release instruction to a first storage resource allocator of low-priority tasks; and

send a video memory resource allocation instruction to a second storage resource allocator of high-priority tasks;

a first storage resource allocator configured to release at least a part of video memory resources occupied by the low-priority task according to the release instruction; and

a second storage resource allocator configured to allocate the video memory resources to the high-priority task according to the allocation instruction, so as to execute the high-priority task at least according to tensor data in a video memory space.

[Claim 19] The system according to claim 18, wherein the allocator is further configured to send usage status information of the video memory resources of the tasks; and

the coordinator is further configured to send the video memory resource release instruction to the allocator if the information satisfies a video memory resource release condition.

[Claim 20] The system according to claim 19, wherein the allocator is further configured to send the information to the coordinator according to a preset period.

[Claim 21] A machine learning system, comprising:

a client configured to send priority information of machine learning tasks to a server; and

a server configured to determine priorities of the plurality of machine learning tasks executed by a graphics processing unit; if video memory resources are to be allocated to a high-priority task, and an amount of allocatable video memory resources is smaller than an amount of video memory resources required by the high-priority task, release at least a part of video memory resources occupied by a low-priority task; and allocate the video memory resources to the high-priority task, so as to execute the high-priority task at least according to tensor data in a video memory space.