

# Learnware: small models do big

Zhi-Hua ZHOU\* & Zhi-Hao TAN

*National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China*

Received 22 April 2023/Accepted 27 June 2023/Published online 20 October 2023

**Abstract** There are complaints about current machine learning techniques such as the requirement of a huge amount of training data and proficient training skills, the difficulty of continual learning, the risk of catastrophic forgetting, and the leaking of data privacy/proprietary. Most research efforts have been focusing on one of those concerned issues separately, paying less attention to the fact that most issues are entangled in practice. The prevailing big model paradigm, which has achieved impressive results in natural language processing and computer vision applications, has not yet addressed those issues, whereas becoming a serious source of carbon emissions. This article offers an overview of the learnware paradigm, which attempts to enable users not to need to build machine learning models from scratch, with the hope of reusing small models to do things even beyond their original purposes, where the key ingredient is the specification which enables a trained model to be adequately identified to reuse according to the requirement of future users who know nothing about the model in advance.

**Keywords** artificial intelligence, machine learning, learnware

## 1 Introduction

Machine learning has achieved great success, while there are lots of complaints about the requirement of a huge amount of training data (particularly data with labels), the difficulty of adapting a trained model to changing environments, and the embarrassment of catastrophic forgetting when refining a trained model incrementally is demanded, etc. There are great efforts such as weakly supervised learning [1] trying to reduce the requirement of labeled training data, open-environment machine learning [2] trying to enable learning models to adapt to environments, and continual learning [3] trying to help deep neural networks resist forgetting; however, these issues are still far from solved.

Indeed, most efforts have been focusing on one of those concerned issues separately, paying less attention to the fact that most issues are entangled in practice. For example, a well-studied technique of weakly supervised learning for reducing the requirement of labeled training data is to collect and exploit a huge amount of unlabeled data drawn from the distribution the same as that of the labeled training data, paying less attention to the fact that in changing environments the data distributions are subject to change inherently. For another example, an effective approach to cope with changing environments is to emphasize data received in very recent timeslots since the changes have not yet caused significant differences, paying less attention to the fact that the emphasis on very recent data may tend to aggravate the severity of catastrophic forgetting.

There are many other issues, e.g., most ordinary users can hardly produce well-performed models starting from scratch, due to the lack of proficient training skills; in many real-world tasks the data privacy/proprietary issue may disable data sharing, leading to the difficulty of sharing experience among different users; in really big data applications, it is generally unaffordable or even infeasible to hold the whole data to support many passes of scanning.

The prevailing deep learning big model paradigm, which has achieved impressive results in natural language processing and computer vision applications [4,5], has not yet addressed the above issues. Note that each big model is targeted to a task (or task class) planned in advance, generally helpless to others, e.g., a big model trained for face recognition can hardly be helpful to financial futures trading. It would

\* Corresponding author (email: [zhouzh@nju.edu.cn](mailto:zhouzh@nju.edu.cn))

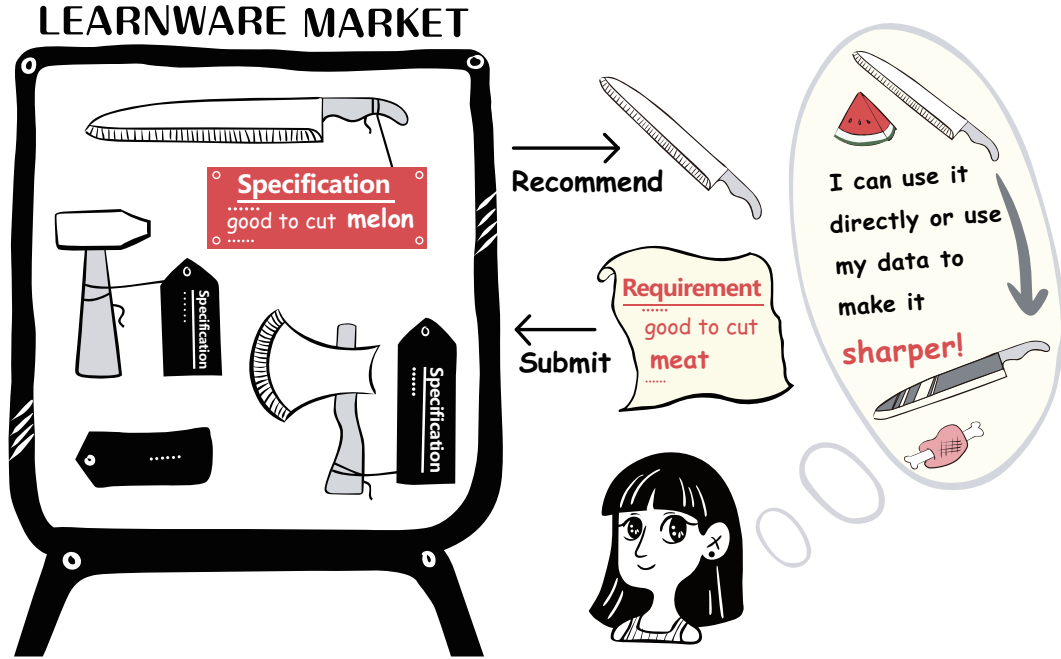


Figure 1 (Color online) An analogy of learnware.

be too ambitious to build a pre-trained big model for every possible task, because the number of possible tasks can be unimaginably big or even infinite. In addition, sadly, the training of big models is becoming a serious source of carbon emissions threatening our environment.

Admitting the usefulness of big models in their specifically targeted tasks, is there any paradigm offering the possibility of tackling the above issues simultaneously?

This article overviews the progress of learnware, a paradigm offering a promising answer to the above question. It attempts to systematically reuse small models to do things that may even be beyond their original purposes, and enables users not to need to build their machine learning models from scratch.

## 2 The learnware proposal

The learnware paradigm was proposed in [6]. A learnware is a well-performed trained machine learning model with a specification which enables it to be adequately identified to reuse according to the requirement of future users who know nothing about the learnware in advance.

The developer or owner<sup>1)</sup> of a trained machine learning model (no matter whether the model is a deep neural network, a support vector machine, or a decision tree) can spontaneously submit her trained model into a learnware market. If the learnware market decides to accept the model, it assigns a specification to the model and accommodates it in the market. The learnware market should not be small, otherwise it can hardly offer help for various tasks; it would be common to accommodate thousands or millions of well-performed models submitted by different developers, on different tasks, using different data, optimizing different objectives, etc.

Once the learnware market has been built, when a user is going to tackle a machine learning task, she can do it in the following way rather than building her model from scratch. As the comic in Figure 1 illustrates, she can submit her requirement to the learnware market, and then the market will identify and deploy some helpful learnware(s) by considering the learnware specification. The learnware can be applied by the user directly, or adapted/polished by user's own data for better usage, or exploited in other ways to help improve the model built from the user's own data. No matter which mechanism for model reuse is adopted, the whole process can be much less expensive and more efficient than building a model from scratch by herself.

The learnware proposal offers the possibility of addressing most issues concerned in Section 1.

1) There are situations where the developer and owner of a trained machine learning model are different. Here, for simplicity, we do not distinguish them and assume that the developer holds all rights of the model.

**Lack of training data.** Strong machine learning models can be attained even for tasks with small data, because the models are built upon well-performed learnwares, and only a small amount of data is needed for adaptation or refinement for most cases.

**Lack of training skills.** Strong machine learning models can be attained even for ordinary users with little training skills, because the users can get help from well-performed learnwares rather than building a model from scratch by themselves.

**Catastrophic forgetting.** A learnware will always be accommodated in the learnware market once it is accepted, unless every aspect of its function can be replaced by other learnwares. Thus, the old knowledge in the learnware market is always held. Nothing to be forgotten.

**Continual learning.** The learnware market naturally realizes continual and lifelong learning, because with the constant submissions of well-performed learnwares trained from diverse tasks, the knowledge held in the learnware market is being continually enriched.

**Data privacy/proprietary.** The developers only submit their models without sharing their own data, and thus, the data privacy/proprietary can be well preserved. Although one could not deny the possibility of reverse engineering the models, the risk would be too small compared with many other privacy-preserving solutions.

**Unplanned tasks.** The learnware market is to be open to all legal developers. Thus, there would exist helpful learnwares in the market unless a task is new to all legal developers. Moreover, some new tasks, though no developer has built models for them specially, could be addressed by selecting and assembling some existing learners.

**Carbon emission.** Assembling small models may offer good-enough performance for most applications; thus, one may have less interest in training too many big models. The possibility of reusing other developers' models can help reduce repetitive development. Besides, a not-so-good model for one user may be very helpful for another user. No training cost wasted.

Though the learnware proposal shows a bright future, there is much work to be done to make it a reality. Sections 3–5 will introduce our progress.

### 3 The design

There are three important entities: developers, users, and the market. The developers are usually machine learning experts who produce and want to share/sell their well-performed trained machine learning models. The users need machine learning services but usually have only limited data and lack machine learning knowledge and skills. The learnware market accepts/buys well-performed trained models from developers, accommodates them in the market, and provides/sells services to users via identifying and reusing learnwares to help users tackle their present tasks<sup>2)</sup>. The basic operation can be decomposed into two stages, as illustrated in Figure 2.

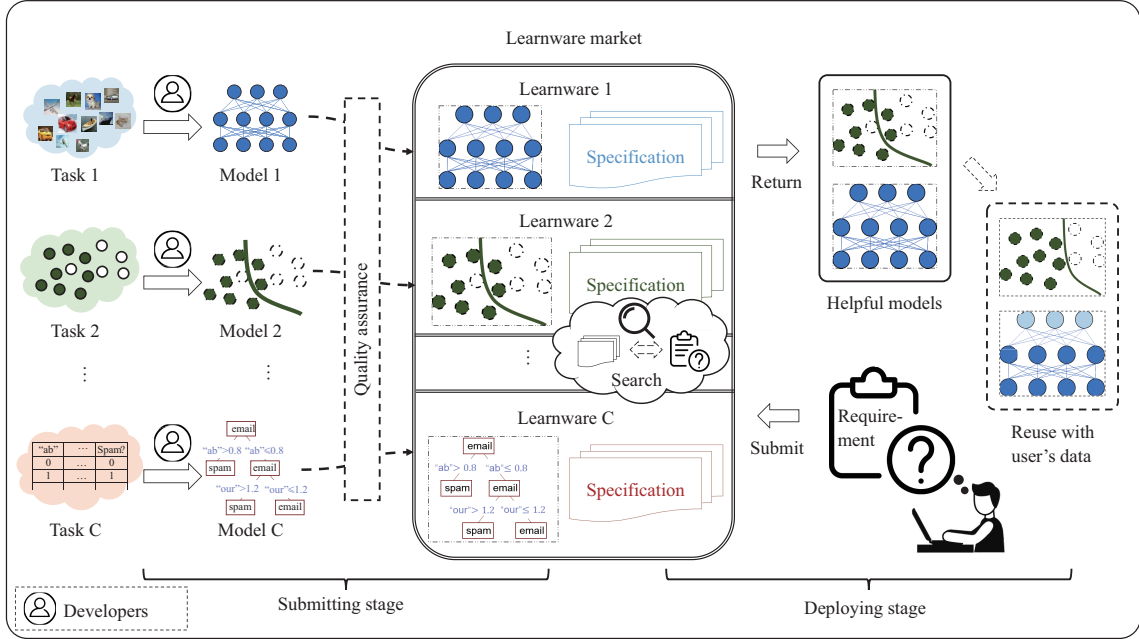
#### 3.1 Submitting stage

In the submitting stage, developers can spontaneously submit their trained models to the learnware market. The market will execute some quality assurance mechanisms, e.g., performance validation, to decide whether a submitted model can be accepted or not. Considering a learnware market which has accommodated millions of models, how to identify potentially helpful models for a new user?

It is evidently undesired to request the user to submit her own data to the market for trials with the models, since this would be too tedious and costly, and more seriously, this could leak user's own data. It is also impossible to utilize straightforward ideas such as “measuring the similarity between the user data and the original training data of models”, as the learnware proposal considers the fact that neither developers nor users would like to leak their own data due to privacy/proprietary issues (it would be easier if their data are free to the market). Thus, our design is based on the constraint that the learnware market has access to neither the original training data of developers nor the original data of users. Besides, it is assumed that users know little about what models have been accommodated in the market.

---

<sup>2)</sup> The learnware proposal implies some possible business relation among the three entities. The user who receives valuable services pays to the market, while the market pays to developers according to the usage of their submitted learnwares. However, the business model is beyond the scope of this article.




**Figure 2** (Color online) The learnware market and the involved two stages.

The key of our solution lies in the specification, which is the core of the learnware proposal. Once the learnware market decides to accept a submitted model, it will assign to the model a specification, which conveys the specialty and utility of the model in some format, without leaking its original training data. For simplicity, consider models corresponding to functions realizing mappings from the input domain  $\mathcal{X}$ , to the output domain  $\mathcal{Y}$ , with regard to the objective  $\text{obj}$ ; in other words, those models reside in a functional space  $\mathcal{F} : \mathcal{X} \mapsto \mathcal{Y}$  w.r.t.  $\text{obj}$ . Each model has a specification. All specifications form a specification space where those of models that are helpful for the same tasks are nearby.

In a learnware market, there will exist heterogeneous models with different  $\mathcal{X}$ , and/or different  $\mathcal{Y}$ , and/or different  $\text{obj}$ . If we call the specification space covering all possible models in all possible functional spaces as the specification world analogically, then each specification space corresponding to one possible functional space can be called a specification island. Designing an elegant specification format covering the whole specification world and enabling all possible models to be efficiently and adequately identified are a grand challenge. Currently, we employ a practical design as follows. The specification of each learnware consists of two parts, where the first part explains which specification island the learnware locates, while the second part, to be introduced in Subsection 3.3, discloses at which location it resides in this island.

The first part can be realized by a string, consisting of a set of descriptions/tags given by the learnware market, about the task, input, output, and objective. Then, according to the descriptions/tags provided in the user requirement, the corresponding specification island can be efficiently and accurately located. Generally, the designer of the learnware market can compose a set of initial descriptions/tags, and the set can grow when the market accepts some new models that could not be accommodated in existing functional spaces, resulting in the creation of new functional spaces and their corresponding specification islands. The learnware market can be ever-increasing as long as its host resource allows.

The specification islands can merge into a larger one, as illustrated in Figure 3. Initially there are two islands, corresponding to functional spaces  $F_1 : \mathcal{X}_1 \mapsto \mathcal{Y}$  w.r.t.  $\text{obj}$  and  $F_2 : \mathcal{X}_2 \mapsto \mathcal{Y}$  w.r.t.  $\text{obj}$ , respectively. When a new model about  $F : \mathcal{X}_1 \cup \mathcal{X}_2 \mapsto \mathcal{Y}$  w.r.t.  $\text{obj}$  is accepted by the learnware market, these two islands can be merged. For example, suppose there are some models on text data (i.e., on  $\mathcal{X}_1$ ) and image data (i.e., on  $\mathcal{X}_2$ ), respectively; once some multi-modal models involving both texts and images are accepted, these text-only and image-only models can become helpful to each other, and they appear to reside in the same extended functional space (in other words, these text-only and image-only models will be kept along with the new multi-modal models). Note that though the learnware market does not have access to the original training data of models, this is still possible because the market can have synthetic data by randomly generating some inputs and feeding them to models, and then



**Figure 3** (Color online) The learnware specification world.

concatenating each input with its corresponding output to construct a data set reflecting the function of a model. In principle, specification islands can be merged if there are common ingredients in  $\mathcal{X}$ ,  $\mathcal{Y}$ , and  $\text{obj}$ . One can imagine when all possible tasks are present; all the specification islands become connected to a non-fragmented unified specification world.

### 3.2 Deploying stage

In the deploying stage, the user submits her requirement to the learnware market, and then the market will identify and return some helpful learnwares to the user. There are two issues, i.e., how to identify learnwares matching the user requirement, and how to reuse the returned learnwares.

The learnware market can accommodate thousands or millions of models. Different to previous machine learning studies about model reuse [7, 8] or domain adaptation [9] where all pre-trained models are assumed to be helpful, there may be only a tiny portion of learnwares helpful for the current user task. Different from multi-task learning [10] where data of the multiple tasks are available in training, and domain-agnostic learning [11] where labeled data of the source domain are available, the learnware market does not assume to have those information. Indeed, efficiently identifying helpful learnwares is quite challenging, particularly when considering the fact that the learnware market has access to neither the original training data of learnwares nor the original data of current user.

With the specification design mentioned in Subsection 3.1, the learnware market can request the user to describe her intention using the set of descriptions/tags, through a user interface or a kind of learnware description language to be designed in the future. Based on such information, the task reduces to how to identify some helpful learnwares in a specification island. The learnware market can provide several anchor learnwares, such as prototypes of functionally similar learnware clusters [12] in the functional space corresponding to the specification island, request the user to test them and return some information, and then identify potentially helpful learners based on this information, as to be explained in Subsection 3.3. The efficiency and scalability of the learnware market can be further improved by maintaining a specification index facilitated with relevant techniques such as hashing. Flexible specification index helps enable the learnware market to be evolvable, such that the market can accommodate an ever-increasing number of learnwares and achieve an increasingly accurate model characterization and identification, sometimes even ascertain the capabilities of models beyond their original purposes.

Once some helpful learnwares are identified and delivered to the user, they can be reused in various ways. The most straightforward one is to apply the received learnware to user's own data directly; if multiple learnwares are received, they can be used to comprise an ensemble [13] for even better performance. The user can also adapt/polish the received learnware(s) by generating a model from her own training



data and putting it to use together with the learnwares(s). Another possible usage is to regard each received learnware as a feature augmentor, by feeding user's data to the learnware, taking its output for each instance as an augmented feature, and then utilizing the augmented data to build the final model.

Note that some helpful learnwares may be trained from tasks that are not exactly the same as the user's current task. For example, there are cases where learnwares with different objectives can be reused to help user's current task, such as that a model which optimizes accuracy can be reused to help a task optimizing AUC, by augmenting a  $\delta$  function based on user's own data [14]. There are also cases where there is no single learnware that can tackle user's task as a whole, but there are multiple learnwares each can tackle a part of user's task separately. In such cases, user task can be tackled in a divide-and-conquer way, as proposed in reusable ensemble [13, pp.184]. Besides, if a very small set of potentially helpful learnwares have been returned, it is possible to reuse them collectively through measuring the utility of each model on each testing instance [15]. Sometimes users may find it difficult to express their requirement accurately. In such cases, it is appealing to reuse the received learnware(s) by adapting/polishing them directly using user's own data. There are preliminary studies that might be somewhat helpful for this purpose, e.g., [7, 8, 14, 16]. Moreover, in learnware deploying it would be beneficial to leverage diverse variants and hardware resource exploitation mechanisms to improve cost efficiency like in some recent explorations [17].

### 3.3 Learnware specification

The learnware specification, ideally, should express/encode important information about every model accommodated in the learnware market to enable them to be identified efficiently and adequately for future users. As mentioned in Subsection 3.1, our current specification design consists of two parts. The first part is a string of descriptions/tags given by the learnware market, based on information submitted by developers, aiming to locate the specification island in which a model resides. Different learnware market enterprises may employ different descriptions/tags.

The second part of the specification plays a crucial role in locating the appropriate place in the functional space  $F: \mathcal{X} \mapsto \mathcal{Y}$  w.r.t.  $\text{obj}$  for the model. Our recent effort is the RKME (reduced kernel mean embedding) specification, based on techniques of a reduced set of KME (kernel mean embedding) [18, 19]. The KME is a powerful technique to map a probability distribution to a point in RKHS (reproducing kernel Hilbert space), whereas the reduced set reserves the ability with a concise representation which does not expose the original data.

Suppose a developer is to submit a model trained from data set  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$ ,  $\mathbf{x}_i \in \mathcal{X}$ ,  $\mathbf{y}_i \in \mathcal{Y}$ . Once the model is trained, the  $\mathbf{x}_i$ s can be fed to the model to get corresponding output  $\hat{\mathbf{y}}_i$ . Note that  $\hat{\mathbf{y}}_i$  is the output of the model instead of ground-truth, and thus, the data set  $\{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^m$  encodes the function of the model; in other words, it offers a function representation of the model. Note that  $\mathbf{x}_i$ s in addition to  $\mathbf{x}_i, i \in [1, m]$  can also be generated and fed to the model for a more thorough representation. This idea has been explored in [20] for learning a simpler model with comparable or even better performance from an original complicated model; a similar idea later has been called knowledge distillation [21]. Here, we take it as the basis for constructing the first part of the RKME specification. For simplicity, let  $\mathbf{z}_i$  denote  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$ , and the function of the model is encoded in the distribution of  $\mathbf{z}_i$ . Then, the market will generate the reduced set representation by minimizing the distance measured by the RKHS norm [22], as

$$\min_{\beta, \mathbf{t}} \left\| \frac{1}{m} \sum_{i=1}^m k(\mathbf{z}_i, \cdot) - \sum_{j=1}^n \beta_j k(\mathbf{t}_j, \cdot) \right\|_{\mathcal{H}}^2, \quad (1)$$

where  $k(\cdot, \cdot)$  is the kernel function corresponding to the RKHS  $\mathcal{H}$ ,  $n \ll m$ , both decided by the learnware market and given to the developer. The solved  $(\beta, \mathbf{t})$ , which offers a much more concise representation very different from the original data  $\mathbf{z}$ , will be submitted by the developer for the second part of the model specification.

In the deploying stage, if the user has many training data, the market can help her construct the RKME requirement to submit. Then, by matching the RKME specifications with the user requirement, the market can identify and return the learnware with the smallest distance in the RKHS norm. The market can also identify multiple helpful learnwares whose weighted combination of RKME specifications has the smallest distance to the user requirement. If the user does not have sufficient training data for constructing an RKME requirement, the learnware market can send several anchor learnwares to the

user. By feeding her own data to these anchor learnwares, some information such as (precision, recall) or other performance indicators, can be generated and returned to the market. These information could help the market identify potentially helpful models, e.g., by identifying models that are far from anchors exhibiting poor performance whereas close to anchors exhibiting relatively better performance in the specification island<sup>3)</sup>.

Note that in the procedures described above, neither the training data of developers nor that of users are leaked to the learnware market.

## 4 Some theoretical results

The RKME specification is based on RKME  $\tilde{\Phi}$ , which aims to make a good representation by constructing a reduced set to approximate the empirical KME  $\Phi = \int_{\mathcal{X}} k(\mathbf{x}, \cdot) dP(\mathbf{x})$  of the underlying distribution. Theoretically, when the kernel function satisfies  $k(\mathbf{x}, \mathbf{x}) \leq 1$  for all  $\mathbf{x} \in \mathcal{X}$ , with probability at least  $1 - \delta$ , we have the guarantee that [22–24]

$$\|\tilde{\Phi} - \Phi\|_{\mathcal{H}} \leq 2\sqrt{\frac{2}{n}} + \sqrt{\frac{1}{m}} + \sqrt{\frac{2 \log(1/\delta)}{m}}, \quad (2)$$

where  $n, m$  are the size of the RKME reduced set and the original data, respectively. It is known that when using characteristic kernels such as the Gaussian kernel, KME can capture all information about the distribution [25]. Besides, when the RKHS of the kernel function is finite-dimensional, RKME enjoys a linear convergence rate  $O(e^{-n})$  to empirical KME [26]; even for infinite-dimensional RKHS, it has been proven constructively that RKME can enjoy  $O(\sqrt{d}/n)$  convergence rate under  $L_{\infty}$  measure, where  $d$  is the dimension of original data [27, 28]. Therefore, the RKME is guaranteed to be a good estimation of KME and a valid representation for data distribution that encodes the ability of a trained model.

The risk on the user task can be bounded under some assumptions, such as the assumption that the distribution corresponding to the task of user matches that of a learnware, or the assumption that it can be approximated by a mixture of distributions corresponding to a set of learnwares' tasks, i.e.,

$$\mathcal{D}_u = \sum_{i=1}^N w_i \mathcal{D}_i, \quad (3)$$

where  $\mathcal{D}_u$  is the distribution corresponding to user task,  $N$  is the number of learnwares, and  $\mathcal{D}_i$  are their corresponding distributions,  $\sum_{i=1}^N w_i = 1$  and  $w_i \geq 0$ . These two assumptions are called task-recurrent and instance-recurrent assumptions, respectively [22]. Besides, assume that all learnwares are well-performed ones, i.e.,

$$\mathbb{E}_{\mathcal{D}_i} [\ell(\hat{f}_i(\mathbf{x}), \mathbf{y})] \leq \epsilon, \quad \forall i \in [N], \quad (4)$$

where  $\hat{f}_i$  is the function corresponding to the  $i$ -th learnware,  $\ell$  is the loss function,  $\mathbf{y}$  is assumed to be determined by a ground-truth global function  $h$ .

Under these assumptions, recent studies have attempted to bound the risk on user task [22, 23]. Consider the task-recurrent assumption and select the learnware  $(\hat{f}_i, \tilde{\Phi}_i)$  with the smallest RKHS distance  $\eta$  according to RKME, given the loss function

$$|\ell(\hat{f}_i(\mathbf{x}), h(\mathbf{x}))| \leq U, \quad \forall \mathbf{x} \in \mathcal{X}, \quad \forall i \in [N]. \quad (5)$$

We have the following result for empirical risk on user task:

$$\hat{\mathbb{E}}_{\mathcal{D}_u} [\ell(\hat{f}_i(\mathbf{x}), \mathbf{y})] \leq \epsilon + U\eta + O\left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}}\right). \quad (6)$$

As for the instance-recurrent assumption and the 0/1-loss,

$$\ell_{01}(f(\mathbf{x}), \mathbf{y}) = \mathbb{I}(f(\mathbf{x}) \neq \mathbf{y}), \quad (7)$$

---

<sup>3)</sup> Both above processes can be realized with a user interface at user's side, without leaking user's training data to the learnware market.

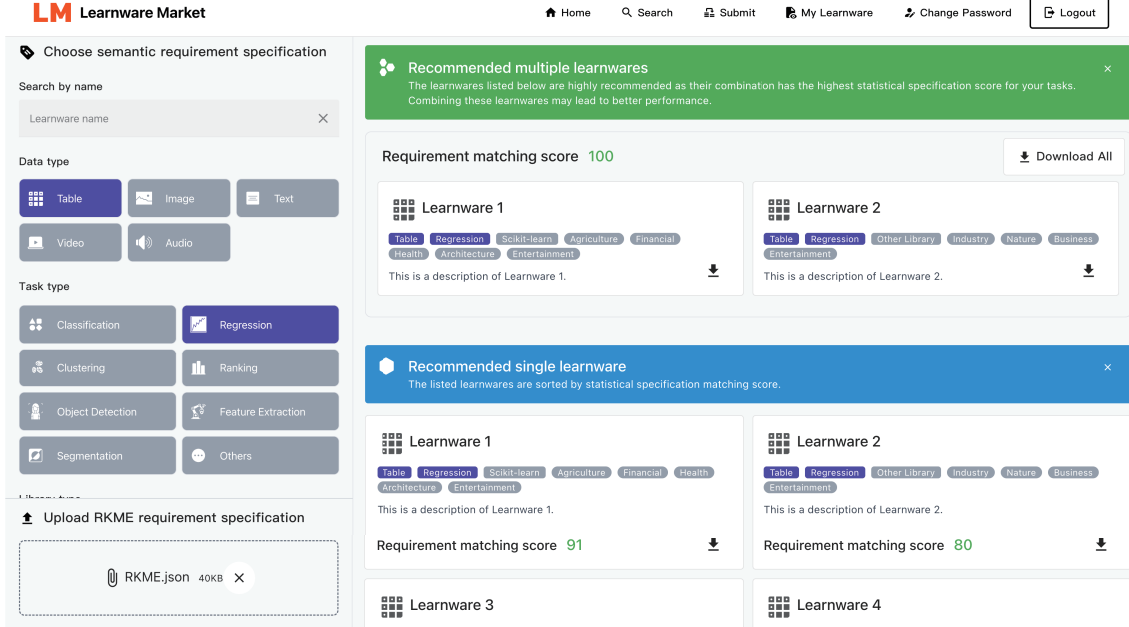


Figure 4 (Color online) Interface of a simple prototype.

a more general result of generalization error on user task has been attained [23],

$$\mathbb{E}_{\mathcal{D}_u} [\ell_{01}(f(\mathbf{x}), \mathbf{y})] \leq \epsilon + R(g), \quad (8)$$

where  $R(g) = \sum_{i=1}^N w_i \mathbb{E}_{\mathcal{D}_i} [\ell_{01}(g(\mathbf{x}), i)]$  represents the weighted risk of any learnware selector  $g(\mathbf{x})$ , which takes unlabeled data as input and assigns it to the proper model, and  $f(\mathbf{x}) = \hat{f}_{g(\mathbf{x})}(\mathbf{x})$  is the final model for user task.

There are efforts trying to enable the learnware market to handle unseen jobs [23], where the user task involves some unseen parts that have never been handled by current learnwares in the market, and a more general theoretical analysis is presented based on mixture proportion estimation [29,30].

There are studies in the field of ensemble learning [13] showing that assembling a group of weak machine learning models can reach strong performance; e.g., it has been proven that weak learnability equals to strong learnability [31], implying that small weak models can be boosted into a strong model if adequately assembled. For a learnware market, however, it is not wise to assume that for arbitrary specific applications there must exist in this market some models that can be exploited to offer strong performance. In practice, the anchor learnwares can offer an indicator, e.g., if the performance of all anchor learnwares is very poor then it may suggest that there is little hope to find helpful learnwares in this market, while theoretical study remains to be explored.


## 5 A simple prototype

A simple prototype learnware market has been implemented for experiments<sup>4)</sup>, with an interface shown in Figure 4. The left-hand panel is for the user to submit requirement specification (including semantic part and/or RKME part), while the right-hand panel returns learnwares identified from the market, with requirement matching scores showing how well the learnwares match the user requirement estimated via RKME specification.

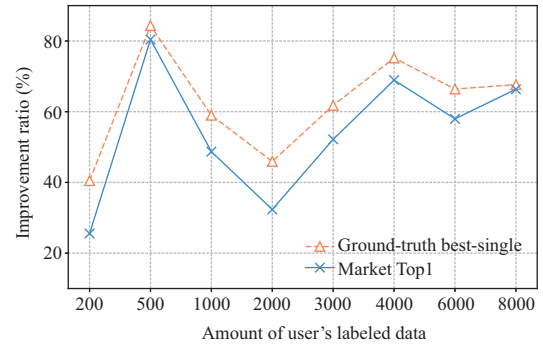
The market accommodates 53 models about sales forecasting. They are with different model types and trained from different data sets, though the input space, output space, and objectives are the same. Thus, the specifications of these models reside in the same specification island realized with the RKME specification. Experiments are conducted to simulate the scenario where a new user, who plans to build her own sales forecasting model, is to get help from the learnware market. The following approaches are tried: to use the best-single model identified from the market directly, to use the best-three models

4) An open-source experimental system with documentation will be made available soon.





**Figure 5** (Color online) Sales forecasting task. Benefit from learnware market in contrast to building a model from scratch using user's own data.



**Figure 6** (Color online) Performance difference of the best-single model (ground-truth) and the Top1 model (identified according to RKME specification matching) in contrast to average model (five random runs).

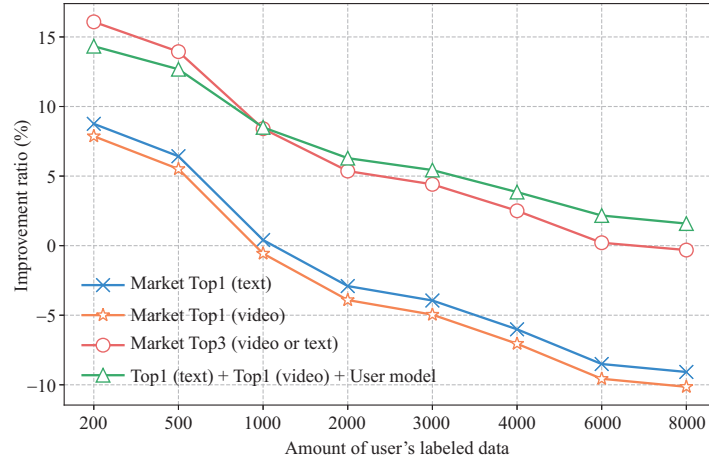
returned from the market via ensemble averaging, to use the best-two models returned from the market together with a model trained from user's own data via ensemble averaging, respectively. The models returned from the learnware market are identified based on RKME specification matching. Figure 5 shows the performance improvement ratio of those approaches against that of the model trained by using user's own data only. Note that the user model is developed with rich machine learning expertise, and when the user has 8000 labeled data, the user model performance is highly competitive with the best model in learnware market.

Figure 5 exhibits that by resorting to the learnware market, the user can get much better models than simply building a model from scratch by using her own data, especially when she has only a small amount of labeled data. In particular, when the user has only 200–1000 labeled data, using the best-single model identified from the learnware market (i.e., Top1) brings more than 20% performance improvement. This verifies that the learnware paradigm can offer a remedy to the lack of training data. It is also noteworthy that when a group of models is identified from the learnware market (i.e., Top3), the improvement against the user's own model is always apparent. Even when the user has 8000 labeled data such that the performance of her own model is highly competitive to the best model and better than the second-best model in the market, these models from learnware market can still be helpful as 10% improvement is observed when they are used together with user's own model (i.e., Top2+User model).

Figure 6 plots the performance of the Top1 model identified according to RKME specification matching in the above experiments, the ground-truth best-single model in the learnware market, against the average model performance from five random runs of selection. It can be seen that the performance of the Top1 model is far superior to that of the average model, and quite close to the ground-truth best-single model in most cases, verifying that the RKME specification matching is effective.

The prototype learnware market also accommodates 6 models about sentiment analysis on video data, and 6 models on textual data with the same output space. A new user, who has some data involving both video and textual information, resorts to the learnware market, and the following approaches are tried to use the best-single video-only or text-only model identified from the market, to use these two models together with a model trained from user's own data via ensemble averaging, to use the best-three models (no matter whether they are video-only or text-only) returned from the market via ensemble averaging, respectively. Figure 7 shows the performance improvement ratio of those approaches against that of the model trained by using user's own data only.

Figure 7 exhibits again that by resorting to the learnware market, the user can get much better results than using her own data only to train a model, especially when she has just a small amount of data. In particular, it can be seen from Figure 7 that when the user has more than 1000 labeled data, the improvement of the best-single models identified from learnware market (i.e., Top1 (text) or Top1 (video)) against the user model is both negative, showing that user's own model is better than the best-single models identified from the learnware market; this is not strange because this is a multi-modal task whereas none of models in the current learnware market was developed for multi-modal tasks. However, it is amazing that when user has less than 1000 labeled data, either the Top1 (text) or Top1 (video) model can do 5% or even 10% better than user model despite the fact that they were not



**Figure 7** (Color online) Sentiment analysis task. Benefit from learnware market in contrast to building a model from scratch using user's own data.

developed for multi-modal tasks whereas the user model was trained exactly for the current multi-modal task. If the three best-single models identified from learnware market (i.e., Top3) are used, though just by the simple ensemble averaging mechanism, performance improvement is visible even when user has 5000 labeled data, despite the fact that none of the Top3 models were trained for multi-modal tasks. These observations verify our argument that learnwares can be useful beyond their original purposes. Furthermore, performance improvement is always visible if user model is employed together with the Top1 video-only and text-only models, implying that it is beneficial to exploit user's own data to adapt/polish the models obtained from learnware market.

It is worth highlighting that such a task, i.e., building a model for sentiment analysis from multi-modal data involving both video and textual information, has never been tackled by previous developers and no model for the exact task exists in the learnware market. This verifies that some new tasks, though no developer has built model for them specifically, can be addressed by selecting and assembling some existing learners. Note that once the user model on multi-modal data is submitted to the learnware market, the two specification islands that correspond to video-only and text-only models, respectively, will merge just like that illustrated in Figure 3.

Though our simple prototype and experiments have exhibited promising aspects of the learnware paradigm, they are still preliminary and thorough empirical study is left for future when large-scale especially enterprise-level learnware markets are available.

## 6 Conclusion and future issues

This article provides a brief overview of progress on learnware, a paradigm that seems promising to tackle many concerns of current machine learning techniques, such as the lack of training data and skills, catastrophic forgetting, continual learning, data privacy/proprietary preserving, unplanned tasks, and carbon emission. It would be great if, in the future, users who plan to build their own machine learning models would look into the learnware market first rather than starting from scratch themselves, just like today's programmers looking for useful codes from Github or other codebases.

There are too many issues for future exploration. First, ideally, the learnware specification should enable well-performed models helpful for the same tasks to locate nearby, whereas our current design is making models with similar functions locate nearby. Considering a user task which can be collectively tackled by several learnwares, one possibility is to tackle the task in a divide-and-conquer way and then look for helpful learnwares for each sub-task. This is to be explored in future. Second, currently the learnwares are assumed to be based on well-performed models whose function can be represented by its training data distribution, whereas in practice the models submitted by developers can be less well-performed. The quality assurance as well as its influence on the identification and reuse procedures is to be studied in future. Third, when the user does not have sufficient data for distribution estimation, as mentioned in Subsection 3.3, some anchor learnwares are to be sent to the user. This can be realized by selecting prototype models through functional space clustering, and more interesting designs are to be

explored in future. Note that our current design tries to assign each model to one location. It is, however, often the case that one model can be helpful for a variety of tasks. To enable one model to be located in multiple suitable specification islands simultaneously is another interesting future issue. Besides, to explore various ways to merge specification islands is also interesting. The learnware market also offers a platform to study the possible “intellectual ability emergence” when models in the market are allowed to have some kind of interaction. Furthermore, though there are some theoretical efforts, it is still far from establishing a thorough theoretical framework for the learnware paradigm.

It is worth emphasizing that the learnware market is a fundamental infrastructure. Though it is possible to be built via volunteering service, in the long run enterprise-level learnware markets may be more favorable. This is because for learnware markets accommodating millions or even more learnwares, many involved issues such as the compressed storage, concurrency control, high throughput, and low latency, need to be considered and the maintenance cost cannot be ignored. It would be beneficial to design credit incentive and versioning control mechanisms to encourage developers to submit upgraded models, and the learnware market can keep multiple versions of learnwares to provide cost-sensitive help to users; for this purpose, meta-data about versions and prices need to be maintained, in addition to the learnware models and specifications stored together in the market. It is expected that some learnware market enterprises may emerge. They build and maintain large-scale learnware markets, try to attract expert developers to submit excellent learners by giving credits such as monetary rewards, and get payment from end-users who get valuable help from the learnware market (the payment must be much less expensive than building a model from scratch by users themselves); this may give birth to a learnware industry.

**Acknowledgements** This work was supported by National Natural Science Foundation of China (Grant No. 62250069).

## References

- 1 Zhou Z-H. A brief introduction to weakly supervised learning. *Natl Sci Rev*, 2018, 5: 44–53
- 2 Zhou Z-H. Open-environment machine learning. *Natl Sci Rev*, 2022, 9: nwac123
- 3 Delange M, Aljundi R, Masana M, et al. A continual learning survey: defying forgetting in classification tasks. *IEEE Trans Pattern Anal Machine Intell*, 2022, 44: 3366–3385
- 4 Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017. 5998–6008
- 5 Brown T, Mann B, Ryder N, et al. Language models are few-shot learners. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020. 1877–1901
- 6 Zhou Z-H. Learnware: on the future of machine learning. *Front Comput Sci*, 2016, 10: 589–590
- 7 Yang Y, Zhan D C, Fan Y, et al. Deep learning for fixed model reuse. In: *Proceedings of Association for the Advancement of Artificial Intelligence*, 2017. 2831–2837
- 8 Zhao P, Cai L W, Zhou Z-H. Handling concept drift via model reuse. *Mach Learn*, 2020, 109: 533–568
- 9 Mansour Y, Mohri M, Rostamizadeh A. Domain adaptation: learning bounds and algorithms. In: *Proceedings of the 22nd Conference on Learning Theory*, 2009
- 10 Zhang Y, Yang Q. A survey on multi-task learning. *IEEE Trans Knowl Data Eng*, 2022, 34: 5586–5609
- 11 Peng X, Huang Z, Sun X, et al. Domain agnostic learning with disentangled representations. In: *Proceedings of International Conference on Machine Learning*, 2019. 5102–5112
- 12 Xie Y, Tan Z H, Jiang Y, et al. Identifying helpful learnwares without examining the whole market. In: *Proceedings of the 26th European Conference on Artificial Intelligence*, 2023
- 13 Zhou Z-H. *Ensemble Methods: Foundations and Algorithms*. New York: Chapman & Hall/CRC, 2012
- 14 Li N, Tsang I W, Zhou Z-H. Efficient optimization of performance measures by classifier adaptation. *IEEE Trans Pattern Anal Mach Intell*, 2013, 35: 1370–1382
- 15 Ding Y-X, Zhou Z-H. Boosting-based reliable model reuse. In: *Proceedings of the 12th Asian Conference on Machine Learning*, 2020. 145–160
- 16 Kuzborskij I, Orabona F. Stability and hypothesis transfer learning. In: *Proceedings of International Conference on Machine Learning*, 2013. 942–950
- 17 Romero F, Li Q, Yadwadkar N J, et al. INFaaS: automated model-less inference serving. In: *Proceedings of USENIX Annual Technical Conference*, 2021. 397–411
- 18 Schölkopf B, Mika S, Burges C J C, et al. Input space versus feature space in kernel-based methods. *IEEE Trans Neural Netw*, 1999, 10: 1000–1017
- 19 Berline A, Thomas-Agnan C. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Berlin: Springer, 2011
- 20 Zhou Z-H, Jiang Y. NeC4.5: neural ensemble based C4.5. *IEEE Trans Knowl Data Eng*, 2004, 16: 770–773
- 21 Hinton G, Vinyals O, Dean J. Distilling the knowledge in a neural network. 2015. ArXiv:1503.02531

- 22 Wu X-Z, Xu W, Liu S, et al. Model reuse with reduced kernel mean embedding specification. *IEEE Trans Knowl Data Eng*, 2023, 35: 699–710
- 23 Zhang Y-J, Yan Y-H, Zhao P, et al. Towards enabling learnware to handle unseen jobs. In: *Proceedings of Association for the Advancement of Artificial Intelligence*, 2021. 10964–10972
- 24 Muandet K, Fukumizu K, Sriperumbudur B, et al. Kernel mean embedding of distributions: a review and beyond. *FNT Machine Learn*, 2017, 10: 1–141
- 25 Sriperumbudur B, Fukumizu K, Lanckriet G. Universality, characteristic kernels and rkhs embedding of measures. *J Machine Learning Res*, 2011, 12: 2389–2410
- 26 Bach F, Lacoste-Julien S, Obozinski G. On the equivalence between herding and conditional gradient algorithms. In: *Proceedings of International Conference on Machine Learning*, 2012. 1355–1362
- 27 Karnin Z, Liberty E. Discrepancy, coresets, and sketches in machine learning. In: *Proceedings of the 32nd Conference on Learning Theory*, 2019. 1975–1993
- 28 Phillips J M, Tai W M. Near-optimal coresets of kernel density estimates. *Discrete Comput Geom*, 2020, 63: 867–887
- 29 Ramaswamy H G, Scott C, Tewari A. Mixture proportion estimation via kernel embeddings of distributions. In: *Proceedings of International Conference on Machine Learning*, 2016. 2052–2060
- 30 du Plessis M C, Niu G, Sugiyama M. Class-prior estimation for learning from positive and unlabeled data. *Mach Learn*, 2017, 106: 463–492
- 31 Schapire R E. The strength of weak learnability. *Mach Learn*, 1990, 5: 197–227