



HAL
open science

Modularity in Deep Learning: A Survey

Haozhe Sun, Isabelle Guyon

► **To cite this version:**

Haozhe Sun, Isabelle Guyon. Modularity in Deep Learning: A Survey. SAI 2023 - Intelligent Computing Conference, Jun 2023, London, United Kingdom. pp.561-595, <10.1007/978-3-031-37963-5_40>. <hal-04224278>

HAL Id: hal-04224278

<https://hal.science/hal-04224278v1>

Submitted on 1 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Modularity in Deep Learning: A Survey

Haozhe Sun ¹ Isabelle Guyon ^{1,2}

¹ LISN/CNRS/INRIA, Université Paris-Saclay, France

² ChaLearn, USA

Email: haozhe.sun@universite-paris-saclay.fr

Abstract

Modularity is a general principle present in many fields. It offers attractive advantages, including, among others, ease of conceptualization, interpretability, scalability, module combinability, and module reusability. The deep learning community has long sought to take inspiration from the modularity principle, either implicitly or explicitly. This interest has been increasing over recent years. We review the notion of modularity in deep learning around three axes: data, task, and model, which characterize the life cycle of deep learning. Data modularity refers to the observation or creation of data groups for various purposes. Task modularity refers to the decomposition of tasks into sub-tasks. Model modularity means that the architecture of a neural network system can be decomposed into identifiable modules. We describe different instantiations of the modularity principle, and we contextualize their advantages in different deep learning sub-fields. Finally, we conclude the paper with a discussion of the definition of modularity and directions for future research.

1 Introduction

Modularity is a general principle present in many fields such as biology [22, 29, 52, 58, 59, 80, 81, 84, 108, 140, 180, 186, 195, 241], complex systems [217, 218], mathematics [14, 31], system design [55, 91, 167, 177, 208], computer science [17, 83], graph theory [93, 168, 170, 182]. While sharing the same name, there is no universally agreed upon definition of modularity [26]. However, we can identify a shared definition [9, 207]: in general, modularity is the property of an entity whereby it can be broken down into a number of sub-entities (referred to as modules). This definition has different instantiations in different fields with their nuances [205] from which various properties may arise. Such field-specific properties include autonomy of modules (limited interaction or limited interdependence between modules) [7, 14, 15, 17, 52, 87, 96, 113, 119, 167, 170, 177, 208, 260], functional specialization of modules [52, 59, 80, 84, 91, 140, 195], reusability of modules [6, 12, 14, 43, 51, 55, 61, 142, 174–176, 184, 188, 192, 207], combinability of modules [6, 12, 144, 151, 165, 177, 184, 188, 239], replaceability of modules [174, 175, 177].

As a general principle, modularity is a descriptive property and an organizational scheme. It is a means of representing entities (data, tasks, models) to be able to manipulate them, conceptually or practically [17, 55, 92, 177]. Though modular entities are not necessarily hierarchical [177], many modular entities have a hierarchical structure [217] in the sense that multiple modules of a lower hierarchy level can form one module of a higher hierarchy level. The modules of the lower hierarchy level are of finer granularity than those of the higher hierarchy level. At the same level of the hierarchy, modules can refer to an exclusive division of the overall entity (hard division) or overlapping parts of the overall entity (soft division). The decomposed modules can be homogeneous (similar modules) or heterogeneous (dissimilar modules).

Back to the very beginning of neural network research in the last century, the community started to be interested in bringing the notion of modularity to neural networks [13, 15, 119, 192], this interest has been revived recently [6, 9, 12, 39, 61, 77, 78, 135, 209, 239]. The publication trend (Figure 1) shows an increasing interest in the modularity principle within deep learning over recent years. This survey investigates the notion of modularity in deep learning around three axes: data, task, and model. The organization of the survey is shown in Figure 2.




Figure 1: **Publication trend of “modular deep learning” from 1990 to 2021.** The ratio of the count of publications containing “modular deep learning” and “modular neural network” among publications containing “deep learning” and “neural network”, indexed by Google Scholar. The horizontal axis is the publication year.




Figure 2: **Organization of this survey.** The first three sections discuss how the modularity principle is instantiated in the three axes: data, task, and model architecture. We then cover other modularity notions for completeness. Finally, we discuss the definition of modularity and directions for future research. The introduction and conclusion are ignored in this figure.

2 Data modularity

Data is an entity used to represent knowledge and information. In the context of machine learning and deep learning, it can take various forms *e.g.*, image, audio sound, and text. Data samples can be interpreted as points in a high dimensional space (fixed-length dense vectors) [8, 136, 138]. A collection of data samples is a dataset. Datasets can be used to train or test deep learning models, referred to as training or test datasets. In these scenarios, data is the input of deep learning models (neural networks) [94].

Data modularity is the observation or creation of data groups; it refers to how a dataset can be divided into different modules for various purposes. The division of the dataset into modules facilitates conception and data manipulation. Data modularization can influence the training of learning machines [71, 100, 227]. Some algorithms leverage data modularity so that each data module is processed by a different solver [187].

We identify two types of data modularity: intrinsic data modularity and imposed data modularity. Intrinsic data modularity means identifiable dataset divisions naturally in data, which a human practitioner does not introduce. Imposed data modularity means identifiable dataset divisions that a human practitioner introduces. The rationale of this taxonomy is that when the dataset reaches the practitioner who analyses it, it already contains some form of intrinsic modularity, including that stemming from the class labels. The people who collect the data are not considered practitioners.

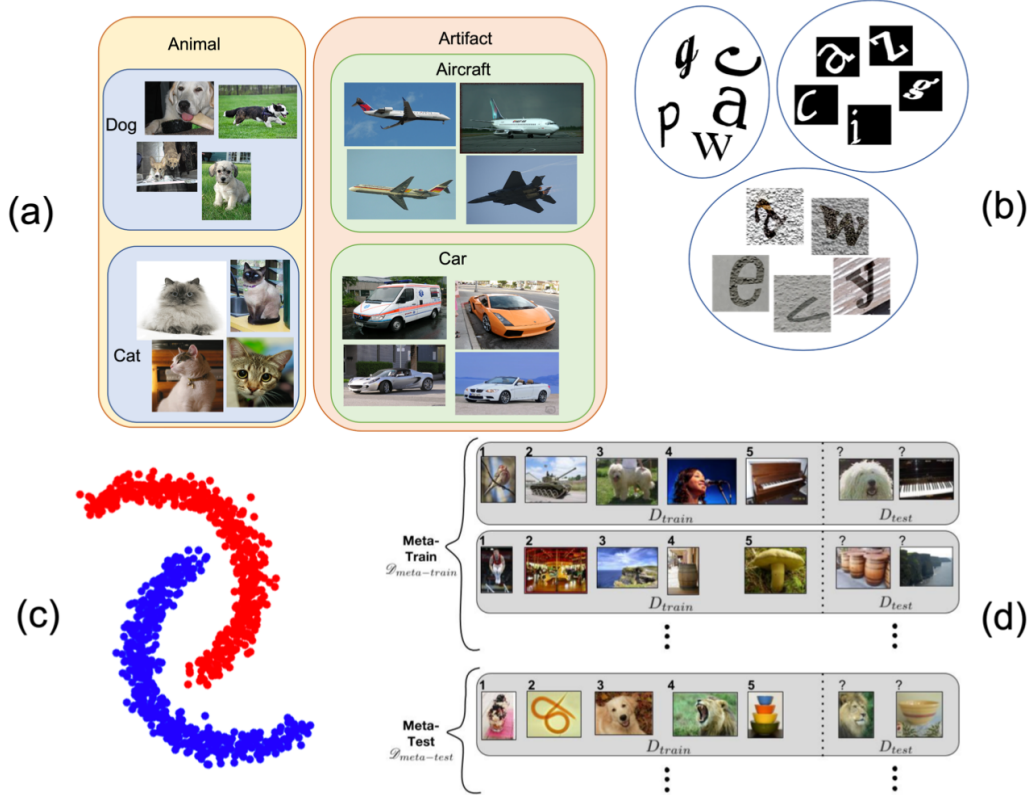


Figure 3: **Illustration of modularity in data.** (a) intrinsic data modularity based on super-classes, images, and class hierarchy in ImageNet [63]; (b) intrinsic data modularity based on styles characterized by a set of metadata, the upper-left circle contains black-on-white characters, the upper-right circle contains white-on-black characters, the lower circle contains characters with natural foreground and background, all characters are drawn from the same set of classes (small-case Latin characters), these three circles illustrate the division of a character dataset based on its metadata; (c) intrinsic manifolds in the form of a moon dataset, where each data manifold can be considered as a module; (d) few-shot learning episodes, reprinted from [191]. (a), (b) and (c) are examples of intrinsic data modularity, (d) is an example of imposed data modularity.

2.1 Intrinsic data modularity

Intrinsic data modularity means identifiable dataset divisions naturally in data, which are not introduced by a human practitioner.

Any supervised learning datasets can be divided according to the classes (labels); data points belonging to the same class are supposed to be close to each other in a hidden space, which allows for solutions of classification algorithms. Classes sharing common semantics can be further grouped to form super-classes. For example, ImageNet [63] has a class hierarchy (see Figure 3 (a)) which is used by Meta-Dataset [235]. Omniglot dataset [142] and OmniPrint datasets [227] contain character images organized in scripts, each script (super-class) contains several characters (classes); Meta-Album dataset [236] is a meta-dataset including 40 datasets, where each dataset can be considered as a super-class. The super-classes provide information about class similarity, allowing splitting datasets according to the semantics [258].

In addition to the classes or super-classes, data points can also be grouped by one or several metadata such as time, location, and gender. Such metadata is available with the Exif data of photos. The OmniPrint data synthesizer generates data together with a comprehensive set of metadata, including font, background, foreground, margin size, shear angle, rotation angle, etc. [227] (see Figure 3 (b)). The NORB dataset collected stereo image pairs of 50 uniform-colored toys under 36 angles, 9 azimuths, and 6 lighting conditions, where the angles, azimuths, and lighting conditions serve as the metadata [145].

Some datasets contain intrinsic clusters in the high-dimensional feature space. Such intrinsic clusters can stem

from the underlying data generative process, where latent categorical variables determine the natural groups of data. An illustrative example is a Gaussian Mixture distribution where data points are assumed to be generated from a mixture of a finite number of Gaussian distributions with unknown parameters [101]. Some datasets have intrinsic manifolds; an illustrative example is the moons dataset as shown in Figure 3 (c), where the two manifolds interlace while preserving an identifiable division, each manifold can be considered as a module. Both of the above examples fall into the category of data clustering. When data samples are interconnected in the form of a graph [154, 249], this is called graph partitioning. One question which arises is how to determine the optimal clustering of a dataset. Luxburg *et al.* [240] argue that there are no optimal domain-independent clustering algorithms and that clustering should always be studied in the context of its end-use.

Multi-modal deep learning aims to build models that can process and relate information from multiple modalities. Here the modality refers to the way in which something happens or is experienced *e.g.*, data in the form of image, text, audio [19]. Multi-modal datasets fall into the category of intrinsic data modularity in the sense that the data in each modality can be considered a module. For example, VQA v2.0 dataset [97] consists of open-ended questions about images; SpeakingFaces dataset [3] consists of aligned thermal and visual spectra image streams of fully-framed faces synchronized with audio recordings of each subject speaking.

2.2 Imposed data modularity

Imposed data modularity means identifiable dataset divisions which are introduced by a human practitioner.

When training deep learning models [94], human practitioners usually divide the whole training dataset into mini-batches, which can be seen as a kind of imposed data modularity. The gradient is computed using one mini-batch of data for each parameter update; one training epoch means passing through all the mini-batches. This iterative learning regime is called stochastic gradient descent [199]. Mini-batches reduce the memory requirement for backpropagation, which makes training large deep learning models possible. On the other hand, batch size also influences learning behavior. Smith *et al.* [221] showed that the benefits of decaying the learning rate could be obtained by instead increasing the training batch size. Keskar *et al.* [131] showed that learning with large batch sizes usually gives worse generalization performance.

Instead of using a sequence of mini-batches sampled uniformly at random from the entire training dataset, curriculum learning [100] uses non-uniform sampling of mini-batches such that the mini-batch sequence exhibits an increasing level of difficulty. A related concept is active learning [193], which assumes that different data points in a dataset have different values for the current model update; it tries to select the data points with the highest value to construct the actual training set.

The model performance is usually tested on few-shot episodes in few-shot learning and meta-learning. Few-shot episodes are typically formed by drawing several classes N from the class pool and several examples K for each selected class, called N -way- K -shot episodes [79, 223] (Figure 3 (d)). For such scenarios, the meta-training phase can employ the same episodic learning regime or not [235], recent studies [141, 242, 244] and competition results [71] suggest that episodic meta-training is not more effective than vanilla pretraining with access to the global class pool.

Data augmentation is a way to generate more training data by applying transformations to existing data [216]. The transformed versions of the same data point can be seen as a module. Some transformations, such as rotation and translation, form a group structure [196]. The effect of such data augmentation can be understood as averaging over the orbits of the group that keeps the data distribution approximately invariant and leads to variance reduction [40].

In addition to splitting the dataset into subsets of samples, each data sample can be split into subdivisions of features, referred to as feature partitioning. A dataset can be represented as a matrix where each row represents one data sample; each column represents one feature dimension. It can then be divided along the sample and feature dimensions. Schmidt *et al.* [207] process each feature partition with a different model. For image classification tasks, input images can be split into small patches that can be processed in parallel [67, 121].

2.3 Conclusion of data modularity

We argue that data without structure contains no useful information for learning dependencies (*e.g.*, between feature and label). Some dependencies boil down to the emergence or the creation of groups. Intrinsic data modularity relates to the semantic relationship between samples and how data samples are similar or dissimilar. Imposed data

modularity, on the other hand, relates to the way that practitioners organize data at hand to better train learning machines.

Future research for data-centric deep learning may investigate the relationship between intrinsic and imposed data modularity. For example, does intrinsic data modularity promote imposed data modularity? How does this interplay affect model training?

Data modularity describes how the input of deep learning models can be modularized. On the other hand, the end goal (the output) of deep learning models can also be modularized, which is the topic of the next section.

3 Task modularity




Figure 4: **Illustration of sub-task decomposition.** The upper figure illustrates the parallel decomposition of a task. The lower figure illustrates the sequential decomposition of a task.

Deep learning models are tools to solve tasks *e.g.*, from the classification of entities to the generation of realistic photos. Solving a task is equal to achieving a corresponding objective. In deep learning, we usually model an objective by an explicit differentiable objective function (also known as a loss function), allowing end-to-end training. This perspective can be generalized to any task, even if the objective function is implicit and does not entail a differentiable form. For example, the task of “purchasing a cup of tea” can be characterized by an indicator function that returns a penalty if no tea can be purchased or a bonus otherwise. In deep learning, tasks are often related to data; but they are different. Given the same dataset, one can define various tasks on top of it. For example, the MNIST dataset can be used either for an image classification benchmarking task [158] or for a pixel sequence classification benchmarking task [96, 139], the OmniPrint-meta[1-5] datasets [227] can be used either for a few-shot learning benchmarking task or for domain adaptation benchmarking task. Tasks define the objective; they are orthogonal to how the end goal should be achieved.

This section presents task modularity *i.e.*, sub-task decomposition. Sub-task decomposition means that a task could be factorized or decomposed into sub-tasks. Sub-task decomposition facilitates conceptualization and problem-solving. The divide-and-conquer principle breaks down a complex problem into easier sub-problems [15, 57, 118, 187]. By solving each individual sub-problem and combining the solutions, the complex problem can be solved more efficiently. The sub-task decomposition facilitates the integration of expert knowledge, and the a priori knowledge can further facilitate problem-solving. Sub-task decomposition can also promote reuse if the overall task is compositional; the solution to sub-tasks may be reused in other tasks [64, 161, 173, 185, 219].

The sub-task decomposition can be categorized into two regimes: parallel decomposition and sequential decomposition (Figure 4). Parallel decomposition means that the sub-tasks can be executed in parallel. Sequential decomposition means that the sub-tasks need to be executed in order; certain sub-tasks cannot be executed before the previous sub-task is finished. In practice, these two regimes can be mixed. For example, a sub-task from a sequential decomposition can be further decomposed parallelly, which leads to a directed acyclic graph workflow.

3.1 Parallel sub-task decomposition

A parallel sub-task decomposition is called homogeneous if the decomposed sub-tasks are similar. One typical example is dividing a multi-class classification problem into multiple smaller classification problems [92]. Given a neural network trained to perform a multi-class classification problem, Csordás *et al.* [61] use parameter masks to

identify subsets of parameters solely responsible for individual classes on their own. Kim *et al.* [133] learn to split a neural network into a tree structure to handle different subsets of classes. They assume that different classes use different features, the tree-structured neural network ensuring that the later layers do not share features across different subsets of classes. Pan *et al.* [174,175] and Kingetsu *et al.* [134] decompose a multi-class classification model into reusable, replaceable and combinable modules, where each module is a binary classifier. Such modules can be recombined without retraining to obtain a new multi-class classifier. These methods can be useful in situations where the classes to be classified frequently change. Abbas *et al.* [2] use transfer learning and class decomposition to improve the performance of medical image classification. Such sub-task decomposability is an implicit prerequisite of the model editing problem [127, 160, 162, 163, 220]. Model editing aims to modify a specific sub-task learned by a trained neural network without damaging model performance on other inputs, *e.g.*, it aims to patch the mistake of the model for a particular sample. If the task cannot be decomposed into disentangled sub-tasks, then model editing cannot be achieved.

A parallel sub-task decomposition is termed heterogeneous if the decomposed sub-tasks are dissimilar; such decomposition is usually problem-dependent and requires expert knowledge of the task at hand. Belay *et al.* [25] decompose the recognition task of Amharic characters into a vowel recognition task and a consonant recognition task to reduce overall task complexity. Cao *et al.* [36] decompose the full self-attention into question-wide and passage-wide self-attentions to speed up inference for question answering tasks. Ding *et al.* [66] decompose the facial recognition task into multiple facial component recognition tasks. Zhou *et al.* [266] decompose the neural network learning task into structure learning and parameter learning to learn equivariance from data automatically. Gatys *et al.* [89] decompose the natural image synthesis task into a content component and a style component, which allows recombining the content and the style in a combinatorial way to generate new images.

3.2 Sequential sub-task decomposition

Sequential sub-task decomposition reflects the sequential pipeline of the task. A simple example is the division of a machine learning task into a preprocessing stage (data cleaning and normalization) and a model inference stage [190].

In reinforcement learning, a complex task can usually be decomposed [219] into a sequence of sub-tasks or steps. An illustrative example is to imagine that the task of manufacturing an artifact Z requires purchasing the raw material X , forging X to produce parts Y , and then assembling the parts Y into the end product Z . Both X and Y can take different values independently ($X \in \{x_1, x_2, x_3, \dots\}, Y \in \{y_1, y_2, y_3, \dots\}$). Different values of X and Y can be recombined, which forms a combinatorial number of possible scenarios to learn. This pipeline can be factorized into three stages: (1) raw material purchase, (2) forging to produce parts, and (3) assembling of parts. Reinforcement learning agents would learn more efficiently if the learning happens at the granularity of the factorized stages instead of the overall task [56]. Furthermore, such a factorization enables the independence of credit assignment [181]; the failure of the overall task can be traced back to the problematic stages, while the other stages can remain untouched. For example, if the raw material is of bad quality, then the purchase sub-task needs to be improved; the forging sub-task and the assembling sub-task do not need to be changed [39].

The sequential pipeline is omnipresent in practical applications *e.g.*, optical character recognition (OCR), natural language processing (NLP). When facing a multi-script (multi-language) recognition task, the pipeline can consist of a script identification stage and a script-specific recognition stage [112, 210], which decouples the domain classifier and the domain-specific solver. The text-in-the-wild recognition task [41] usually consists of decoupled text detector (to localize the bounding box of the text) and recognizer (recognize the text from the bounding box) [41]. Traditional OCR methods also decompose the word recognition task into a character segmentation task and a character recognition task [37, 48, 128, 204]. Traditional NLP pipeline includes sentence segmentation, word tokenization, part-of-speech tagging, lemmatization, filtering stop words, and dependency parsing [125]. In bioinformatics, the scientific workflow (data manipulations and transformations) groups similar or strongly coupled workflow steps into modules to facilitate understanding and reuse [55].

3.3 Conclusion of task modularity

The sub-task decomposition can be parallel, sequential, or mixed (directed acyclic graph). We provided examples from the literature that leverage sub-task decomposition to reduce task complexity or promote the reuse of sub-task solutions. Task modularity can help integrate expert knowledge and promote model interpretability when paired

with model modularity, as will be discussed in the next section.

Future research may focus on how to automate the process of sub-task decomposition or make the problem-dependent sub-task decomposition techniques transferable to other tasks, which is an important step for AutoML. It would reduce the demand for highly qualified deep learning engineers, which can reduce expert bias and entry barriers to deep learning.

4 Model modularity

This section presents model modularity. It means that the architecture of the neural network system (one neural network or a system of neural networks) consists of identifiable sub-entities (modules).

Model modularity is different from task modularity. A task define an objective, task modularity focuses on decomposing the objective into sub-objectives. Model modularity focuses on the architecture of the neural network system, it decomposes the solution into sub-solutions.

4.1 Advantages of model modularity

Model modularity provides ease of conceptual design and implementation. For example, modern neural networks consist of repeated layer/block patterns (modules). Examples include fully-connected neural networks [94], vanilla convolutional neural networks, ResNet [105, 251], Inception [230] and models searched by Neural Architecture Search (NAS) [74, 270]. The design with homogeneous modules allows for a more concise description of the model architecture in the sense of Kolmogorov complexity (short description length) [148, 149]. For example, instead of specifying how each primitive operation (*e.g.*, sum, product, concatenation) interacts in a computational graph, the model can be described as a collection of modules that interact with each other [92]. The standardization of such neural network building blocks (fully-connected layers, convolutional layers) also enabled the development of highly optimized hardware and software ecosystems for fast computation [1, 90, 98, 156, 178].

Together with sub-task decomposition (task modularity), model modularity offers ease of expert knowledge integration [12, 25, 95, 211, 214] and interpretability [118, 135, 137, 184]. Interpretability can have different forms. For example, each neural network module could be assigned a specific interpretable sub-task. On the other hand, selective module evaluation provides insights on how different samples/tasks are related [6, 12, 119, 209] in the context of conditional computation [28].

The model decomposition into modules promotes reusability and knowledge transfer [33]. Though each neural network is typically trained to perform a specific task, its (decomposed) modules could be shared across tasks if appropriate mechanisms promote such reusability. The simplest example would be the classical fine-tuning paradigm of large pretrained models [50, 104, 106, 258]. This paradigm typically freezes the pretrained model and only retrains its last classification layer to adapt it to the downstream task. Pretrained models are typically pretrained on large datasets [201, 225, 254]. The large amount and diversity of training data make pretrained models’ intermediate features reusable for other downstream tasks. More recently, the finer-grained reusability of neural network systems has attracted the attention of researchers. Such methods assume that the tasks share underlying patterns and keep an inventory of reusable modules (each module is a small neural network) [6, 12, 135, 239]. Each module learns different facets (latent factors or atomic skills) of the knowledge required to solve each task. The selective/sparse use and dynamic reassembling/recombination of these modules can promote sample efficiency [184] and combinatorial generalization [6, 12, 62, 117].

Combinatorial generalization is also known as compositional generalization, “infinite use of finite means” [47], and systematic generalization. It aims to generalize to unseen compositions of known functions/factors/words [38, 60, 82, 132, 143, 173, 185], it is the ability to systematically recombine previously learned elements to map new inputs made up from these elements to their correct output [206]. For example, new sentences consist of new compositions of a known set of words. Combinatorial generalization is argued to be important to achieve human-like generalization [23, 114, 115, 134, 142, 146, 153, 164, 183, 184, 237, 243]. Learning different facets of knowledge with different modules in a reusable way could be one solution to combinatorial generalization. Modular systems have been shown effective for combinatorial generalization [197] in various fields *e.g.*, natural language processing [114, 144, 169, 183, 184], visual question answering [12, 16, 62], object recognition [38, 142, 176, 185], and robotics [6, 51, 64, 179].

The modularization of neural network systems promotes knowledge retention. If different knowledge is localized into different modules, targeted knowledge updates and troubleshooting [134, 174, 175] will be possible. This can

alleviate gradient interference of different tasks [126, 155, 261] and catastrophic forgetting [4, 6, 10, 72, 85, 120, 129, 172, 202, 233, 239].

Modular neural network systems facilitate model scaling in two ways. (1) Modular models like fully-connected models and ResNet can be scaled up (or down) by simply stacking more (or less) modules to increase (or decrease) the model capacity to fit larger (or smaller) datasets [105]. (2) Modular methods based on sparsely activated Mixture-of-Experts [209] decouple computation cost from model size. They allow drastically increasing the model capacity without increasing compute cost because only a small fraction of the model is evaluated on each forward pass [21, 49, 68, 75, 102, 209]. The extreme example of these sparsely activated models is Switch Transformer [76] which contains 1.6 trillion parameters, pushing the competition of large model sizes [35, 222] to the next level.

4.2 Typical modules in deep learning models

This section reviews some typical modules in the deep learning literature.

Almost all systems are modular to some degree [205], neural network systems can almost always be decomposed into subsystems (modules) [18] following different points of view. More specifically, they usually consist of a hierarchical structure in which a module of a higher hierarchy level is made of modules of a lower hierarchy level. The elementary layer of modern neural networks (*e.g.*, fully-connected layer, convolutional layer) can be seen as a module on its own. On the other hand, any neural network as a whole can also be considered as a module *e.g.*, in the context of ensemble [268], Mixture-of-Experts [119], and Generative Adversarial Networks (GAN) [95]. Some literature [26, 61, 134, 226] define modules as sub-neural networks where part of the parameters are masked out (set to 0). In these cases, overlapping modules can be obtained when the masks overlap.




Figure 5: **Examples of a module.** (a) a fully-connected layer; (b) a basic ResNet module, reprinted from [105]; (c) an LSTM module, reprinted from [44].

4.2.1 Modules for non-sequential data

Fully-connected layers (Figure 5 (a)) imitate the connections between neurons in biological neural networks but connect every input neuron to every output neuron [94]. In practice, a fully-connected layer is implemented as a matrix multiplication between input data and learnable parameters. Convolutional layers introduce the inductive bias of translation equivariance. Conceptually, a convolutional layer (with a single output channel) can be obtained from a fully-connected layer by enforcing local connectivity and parameter sharing [94]. Local connectivity means that each neuron only connects to a subset of neurons of the previous layer; parameter sharing means that the same learnable parameters are used across receptive fields. In practice, a convolutional layer is implemented as a collection of kernels/filters shifted over the input data [156, 178]. Each kernel performs a dot product between input data and learnable parameters. Depending on the number of dimensions over which kernels are shifted, a convolutional layer is termed *e.g.*, 1D, 2D, 3D. 2D convolutional layers are widely used in computer vision tasks [138, 147]. Locally connected layers are similar to convolutional layers except that they remove the constraint of parameter sharing (across kernels). It helps if one wants to impose local receptive fields while there is no reason to think each local kernel should be the same [94]. Low-rank locally connected layers relax spatial equivariance and provide a trade-off between locally connected layers and convolutional layers. The kernel applied at each position is constructed as a linear combination of a basis set of kernels with spatially varying combining weights. Varying the number of basis kernels allows controlling the degree of relaxation of spatial equivariance [73]. Standard convolutional layers offer translation equivariance; a line of research focuses on generalizing this to other equivariances (rotation,

reflection), referred to as group convolutional layers [24, 53, 54, 65, 88, 246–248]. On the other hand, depthwise separable convolutional layers [46, 109, 213] factorize a standard convolutional layer into a depthwise convolutional layer and a pointwise convolutional layer, which reduces model size and computation.

Multiple layers can be grouped into a building block (a module of a higher hierarchy level). Such examples include the building blocks of ResNet [105], Inception [229, 230], ResNeXt [251], Wide ResNet [262]. Inception [229, 230] has parallel kernels of multiple sizes within each block and merge their results to extract information at varying scales. Inception also includes several techniques to reduce computation cost *e.g.*, factorizing large kernels into smaller kernels and using 1×1 convolution to reduce dimensionality. A ResNet block [105] (Figure 5 (b)) contains a sequence of convolutional layers; it adds a skip-connection (also known as residual connection, identity mapping) from the beginning to the end of the block to alleviate vanishing gradients. Many variants of the ResNet block have been proposed. For example, Wide ResNet [262] increases the block width; ResNeXt [251] aggregates parallel paths within each block.

The block design could be automatically searched instead of handcrafted. In order to narrow down the model search space, some Neural Architecture Search methods [74, 116, 257, 270] automatically search the optimal design pattern for a block (also known as a cell) while fixing the block composition scheme (also known as meta-architecture). Once the block design patterns are searched, the full model is instantiated by repeating the searched blocks following the predefined block composition scheme. For example, NAS-Bench-101 [257] defines the block search space as all possible directed acyclic graphs on V nodes ($V \leq 7$) while limiting the maximum number of edges to 9.

McNeely-White *et al.* [159] report that the features learned by Inception and ResNet are almost linear transformations of each other, even though these two architectures have a remarkable difference in the architectural design philosophy. This result explains why the two architectures usually perform similarly and highlights the importance of training data. This result is corroborated by Bouchacourt *et al.* [30], who argue that invariance generally stems from the data itself rather than from architectural bias.

4.2.2 Modules for sequential data

When the input data is sequential *e.g.*, time series, text, audio, video, Recurrent Neural Networks (RNN) [200] come into play. The RNN module processes the sequential data one at a time; the output (also known as the hidden state) of the RNN module at the previous time step is recursively fed back to the RNN module, which allows it to aggregate information across different time steps. The vanilla RNN module suffers from short-term memory issues; it cannot effectively preserve information over long sequences. To overcome this issue, gated recurrent unit (GRU) [45] and long short-term memory (LSTM) [107] module use gates to control which information should be stored or forgotten in the memory, which allows better preservation of long-term information. In GRU and LSTM modules, gates are neural networks with trainable parameters. While GRU modules are faster to train than LSTM modules, their performance comparison varies depending on the scenario. GRU surpasses LSTM in long text and small dataset scenarios while LSTM outperforms GRU in other scenarios [255].

Contrary to RNN, GRU, and LSTM, which process sequential data one at a time, self-attention layers [238] process the data sequence in parallel. For each data point in a data sequence (*e.g.*, each time step of a time series), a self-attention layer creates three transformed versions, referred to as query vector, key vector, and value vector, through linear transformations. Between each pair of data points, the dot product between the query vector and the key vector of the pair reflects how much those two data points are related within the sequence. These dot products are then normalized and combined with the corresponding value vectors to get the new representation of each data point in the sequence. An enhanced version of self-attention layers is multi-head self-attention layers, which extract different versions of query vector, key vector, and value vector for each data point. Multi-head self-attention layers improve performance by capturing more diverse representations. A transformer block combines multi-head self-attention layers, fully-connected layers, normalization layers, and skip-connections. Models built upon transformer blocks have achieved state-of-the-art performance in a wide range of tasks such as natural language processing [130] and speech synthesis [150]. Transformer models can be applied to image modality by transforming each input image into a sequence of small image patches [67]. Despite the lack of image-specific inductive bias (translation equivariance, locality), vision transformers can achieve state-of-the-art performance when combined with a large amount of training data [20, 67, 103].




Figure 6: **Illustration of module composition.** (a) Sequential concatenation. (b) Ensembling. (c) Tree-structure composition. (d) General Directed Acyclic Graph. (e) Conditional composition. (f) Cooperation composition.

4.3 Composition of modules

Section 4.2 presents typical modules in the literature. Section 4.3 discusses how to organize these modules to form a model (or a module of a higher hierarchy level).

4.3.1 Static composition of modules

Static composition means that the composed structure does not vary with input; the same structure is used for all input samples or tasks.

One straightforward way to compose modules is sequential concatenation (Figure 6 (a)). It implies that multiple (typically homogeneous) modules are sequentially concatenated into a chain to form a model, where a module’s output is the next module’s input. Examples of sequential concatenation include fully-connected models [94] and ResNet models [105]. This composition scheme typically does not assume an explicit sub-task decomposition; the chain of concatenated modules can instead be seen as a series of information extraction steps [5, 234, 258], extracted features transition from low-level to high-level.

Ensembling composition [124, 171, 268], on the other hand, organizes modules in a parallel manner (Figure 6 (b)). The principle of ensembling is to aggregate (*e.g.*, averaging) the results of multiple modules (weaker learners) to obtain a more robust prediction. The rationale is that different modules are expected to provide complementary and diverse views of input data. Each module’s data is processed independently without relying on the other modules at inference time. The regularization method Dropout [224], which randomly deactivates neurons during training, can be seen as an implicit ensemble method of overlapping modules.

Sequential composition and parallel composition can be combined, *e.g.*, in the form of a tree structure (Figure 6 (c)). A typical scenario of tree-structure composition is a model with a shared feature extractor and multiple task-specific heads [215, 265]. All the above composition schemes are special cases of DAG (Directed Acyclic Graph, Figure 6 (d)). The general DAG composition scheme is typically found in models searched by Neural Architecture Search [152, 192, 252].

Cooperation composition (Figure 6 (f)) assumes that each module is a standalone neural network with specific functionality and that these neural networks cooperate during training or inference; it is a neural network system that consists of multiple separate neural networks. Different from ensembling composition, modules in cooperation composition are typically heterogeneous and interact with each other more diversely. For example, siamese networks [34, 42, 122] consists of two neural networks (module) which work together to produce different versions of the input data. Generative Adversarial Networks (GAN) [95, 269] trains a generator under the guidance of a discrim-

inator. The same spirit applies to teacher and student neural networks [231]. Some deep reinforcement learning methods implement the Actor-Critic [228] with two separate new networks, such as AlphaGo [214], A3C [166], ACKTR [250]. Continual learning with deep replay buffer [211] consists of a continual neural network learner and a generative neural network serving as the replay buffer. Some other continual learning methods [10, 202, 233, 239] continuously expanding model capacity for new tasks by adding new modules which work in cooperation with old modules.

4.3.2 Conditional composition of modules

Conditional composition (Figure 6 (e)) is complementary to static composition in the sense that the composed modules are selectively (conditionally, sparsely, or dynamically) activated (used or evaluated) for each particular input. The input conditioning can happen at the granularity of individual sample [12, 119, 135] as well as task [118, 155, 157, 184, 226]. In the literature, this paradigm is also termed conditional computation [27, 28].

The idea of conditional computation can be traced back to Mixture-of-Experts (MoE) introduced in the last century. An MoE is a system composed of multiple separate neural networks (modules), each of which learns to handle a sub-task of the overall task [118, 267] *e.g.*, a subset of the complete training dataset. A gating network computes the probability of assigning each example to each module [119, 123] or a sparse weighted combination of modules [209]. Two issues of MoE are module collapse [135, 164, 209] and shrinking batch size [209], both of which are related to the balance of module utilization. Module collapse means under-utilization of modules or lack of module diversity. Due to the self-reinforcing behavior of the gating network during training, premature modules may be selected and thus trained even more. The gating network may end up converging to always selecting a small subset of modules while the other modules are never used. Shrinking batch size means the batch size is reduced for each conditionally activated module. Large batch sizes are necessary for modern hardware to make efficient inferences because they alleviate the cost of data transfers [209].

MoE can be generalized to *e.g.*, stacked MoE [70, 77, 135, 189, 198] or hierarchical MoE [209, 256] (Figure 7). Eigen *et al.* [70] first explored stacked MoE; they introduced the idea of using multiple MoE with their own gating networks. In order to train stacked MoE, Kirsch *et al.* [135] use generalized Viterbi Expectation-Maximization algorithm, Rosenbaum *et al.* [198] employ a multi-agent reinforcement learning algorithm, Fernando *et al.* [77] use a genetic algorithm. MoE systems do not always have explicit gating networks; for instance, Fernando *et al.* [77] rely on the results of the genetic algorithm to decide the module routing scheme.




Figure 7: **Extension of Mixture-of-Experts (MoE).** (a) A stacked MoE, which stacks multiple MoE layers into a chain. (b) A hierarchical MoE, where a primary gating network chooses a sparse weighted combination of “modules”, each of which is an MoE with its own gating network.

Inspired by MoE, some deep learning methods keep an inventory of reusable specialized modules that can

be conditionally reassembled for each input. This approach has been advocated to promote knowledge transfer, sample efficiency, and generalization. For example, in visual question answering, Neural Module Networks [12, 62, 111] dynamically reassemble modules into a neural network to locate the attention (region of interest) on the questioned image. The question’s parsing guides the reassembling process so that the reassembled model reflects the structure and semantics of the question. For this particular task, the compositionality of modules comes from the compositionality of visual attention. Following the question’s syntax, the reassembled modules sequentially modify the attention onto the questioned image. For example, the module associated with the word “cat” locates the image region containing a cat, and the module associated with the word “above” shifts up the attention. Zhang *et al.* [264] investigated adding new abilities to a generic network by directly transplanting the module corresponding to the new ability, dubbed network transplanting.

Some work relies on the hypothesis that the tasks at hand share some commonalities *i.e.*, hidden factors are shared across tasks. Each hidden factor can be learned by a separate module from the module inventory for transfer learning and meta-learning. For example, Alet *et al.* [6] use simulated annealing to meta-learn an inventory of modules reusable across tasks to achieve combinatorial generalization. The parameters of an inventory of modules are optimized during meta-training; the trained modules are reassembled during the meta-test with an optional parameter fine-tuning process. They demonstrated the utility of their method for robotics tasks. Ponti *et al.* [184] assume that each task is associated with a subset of latent discrete skills from a skill inventory. They try to generalize more systematically to new tasks by disentangling and recombining different facets of knowledge. More precisely, they jointly learn a skill-specific parameter vector for each latent skill and a binary task-skill allocation matrix. For each new task, the new model’s parameter vector is created as the average of the skill-specific parameter vectors corresponding to the skills present in the new task (in addition to a shared base parameter vector).

The conditional composition scheme also has other forms. For example, Teerapittayanon *et al.* [232] save computation on easy input data via early exiting; later layers will be skipped if the intermediate feature’s prediction confidence passes a predefined threshold. Fuengfusin *et al.* [86] train models whose layers can be removed at inference time without significantly reducing the performance to allow adaptive accuracy-latency trade-off. Similarly, Yu *et al.* [259] train models which are executable at customizable widths (the number of channels in a convolutional layer). Xiong *et al.* [253] sparsely activate convolutional kernels within each layer for each particular input sample, which provides an example of the conditional composition of overlapping modules.

4.4 Conclusion of model modularity

Section 4 presents how the notion of modularity is instantiated in the architecture of neural network systems. The structure of neural network modules (Section 4.2) and the way to organize the modules (Section 4.3) provide a complementary view of model modularity.

While all modern neural networks are modular to some extent, different instantiations of the modularity principle offer different advantages (Section 4.1). The advantages include ease of conceptual design and implementation, ease of expert knowledge integration, better interpretability, ease of knowledge transfer and reuse, better generalization and sample efficiency, ease of knowledge retention, ease of troubleshooting, and better scalability.

5 Other notions of modularity

There remain some other notions of modularity in the deep learning literature.

In graph theory, the term “modularity” refers to a measure commonly used in community detection. It measures the density of connections within a community (module) compared to between modules communities [170]. This measure can be applied to graph clustering problems in the form of modularity optimization [32, 110, 203, 212]. Inspired by this measure, Filan *et al.* [78] investigate the parameter clustering pattern that emerged from the training of a neural network. They view a neural network as an undirected weighted graph (edge weights are the absolute value of network parameters) and apply spectral clustering on the obtained graph. They observe that some neural networks trained on image classification tasks have some clustering properties of their parameters: edge weights are stronger within one cluster than between clusters. Watanabe *et al.* [245] have obtained similar results. Béna *et al.* [26] adapted the graph-theoretic modularity measure to define structural modularity and define functional specialization through three heuristic measures. The functional specialization can be intuitively understood as the extent to which a sub-network can do a sub-task independently. To investigate the relationship between structural and functional modularity, they design a scenario where a model with two parallel modules

(with an adjustable number of interconnections) is used to predict whether the parity of the two digits is the same or different. They show that enforcing structural modularity via sparse connectivity between two communicating modules does lead to functional specialization of the modules. However, this phenomenon only happens at extreme levels of sparsity. With even a moderate number of interconnections, the modules become functionally entangled. Mittal *et al.* [164] observed that modular systems (weighted combination of parallel modules) with a good module specialization are good in terms of the overall system performance, however end-to-end training itself is not enough to achieve a good module specialization.

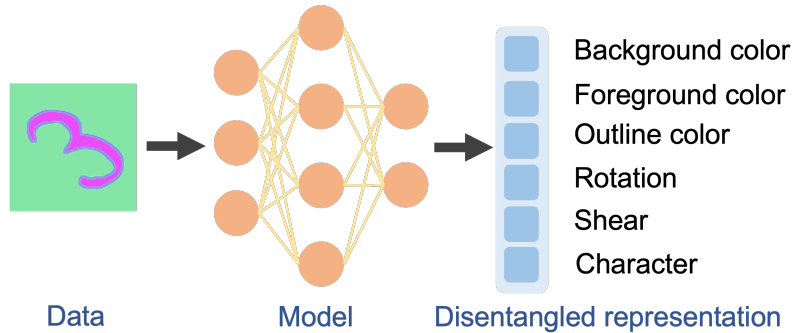


Figure 8: **Illustration of a disentangled representation.**

The term “modularity” is related to the notion of independence in some literature. For example, Galanti *et al.* [87] use modularity to refer to the ability of hypernetworks [99] to learn a different function for each input instance. A line of research has been carried out on learning disentangled representation. Intuitively, disentangled representation aims to reverse the underlying data generating process and retrieve its latent factors into the learned representation (Figure 8). One of the desirable properties of a disentangled representation [69, 194, 263] is “modularity”. In this context, a modular representation is a representation where each dimension of the representation conveys information about at most one latent generative factor.

6 Discussion

Defining modularity is, in itself, a challenging problem. The notion of modularity is present in literature across many different fields [14, 17, 22, 29, 31, 55, 58, 59, 80, 81, 83, 84, 91, 93, 140, 167, 168, 170, 177, 180, 182, 186, 195, 208, 217, 218, 241]. While many researchers have a strong intuition about what it means for an entity to be modular, there has yet to be a universal agreement on what defines modularity. The same is true even within the field of deep learning. As rightly said by Béna *et al.* [26]: “Modularity of neural networks is a bit like the notion of beauty in art: everyone agrees that it’s important, but nobody can say exactly what it means”. We argue that the difficulty of defining modularity stems from the fact that the notion of modularity usually comes with many different properties: replaceability of modules, combinability of modules, reusability of modules, autonomy of modules (limited interaction or limited interdependence between modules), functional specialization of modules. Authors from different fields typically only retain one or two of the above properties to claim an entity to be modular.

In this survey, we define modularity as the property of an entity whereby it can be broken down into a number of sub-entities (referred to as modules). This definition is the prerequisite of the properties mentioned above; it is the greatest common definition of the notion of modularity. By recursively applying the definition of modularity, a modular entity is an entity that can be broken down into sub-entities, where each sub-entity can be further broken down into sub-sub-entities. This recursion can be repeated for discrete entities until the atomic elements (minimum indivisible modules) are reached. In that case, a set of atomic elements $\{a \in D\}$ can formally characterize a discrete entity D ; a subset of atomic elements can then characterize a module $M \subseteq D$. The above framework applies to data modularity (Section 2) and model modularity (Section 4). The reason is that data and models are both discrete: data samples and model parameters are stored in physical computers where everything is represented quantitatively. On the other hand, we need to use a different framework for task modularity because tasks are usually not discrete. As discussed in Section 3, each task can be characterized by an objective function F . In this sense, task modularity can be formally characterized by (objective) function compositions. A task is decomposable if there exists a set of

functions $\{f_1, f_2, \dots\}$ that, when composed together, retrieve the form of the original objective function F .

For discrete entities, one needs to choose the atomic elements. Naively, one could choose each data sample in a dataset and each neuron in a neural network as the atomic elements. However, both choices remain to be discussed because they are indeed not the smallest indivisible modules. Regarding data modularity, the dataset division can happen both at the sample dimension and the feature dimension, which means that each data sample can be divided into smaller elements *e.g.*, feature vectors of reduced length or image patches. Regarding model modularity, the modularization can happen at the granularity of parameters *e.g.*, modules can be obtained by masking out parameters [26, 61, 134, 226]. Consequently, one can choose the scalar numbers stored in physical computers (often represented by floating-point numbers) as the atomic elements. The atomic elements for data are every single dimension of data samples; the atomic elements for models are every single scalar parameters in the neural network. It entails that, in some cases, there needs to be some relationship R among atomic elements $\{a \in D\}$ because any arbitrary subsets of atomic elements do not necessarily form a valid module if the relationship R is broken. In the above example, the relationship R indicates which scalar numbers should come together to form data samples or how to use each scalar parameter along the feedforward computation in the computational graph of neural networks. In consequence, an entity can be a set or a system; a system is a set equipped with relationships R among atomic elements $\{a \in D\}$.

7 Future research

As modularity is a general principle, this survey covered many elements from different sub-fields of deep learning; each sub-field can provide a lot of future avenues of research on its own. To name a few, McNeely-White *et al.* [159] and Bouchacourt *et al.* [30] showed that given the same training data, learned features exhibit similar properties across models with markedly different architectural inductive biases. Is it still worth improving neural network architectures if data dominate learning results? Future research may validate the results of McNeely-White *et al.* [159] and Bouchacourt *et al.* [30] by extending their research to more kinds of models and training datasets in a more systematic way. If these results still hold, one may need to ground these results theoretically. On the other hand, whether neural networks can learn and behave compositionally is still an open question [11, 114]. It entails that we need a domain-agnostic way to test the compositionality of neural networks.

Different aspects of the modularity principle can be further investigated to improve deep learning models. It boils down to designing new deep learning methods that provide *e.g.*, better interpretability, reusability, scalability, and efficiency. While model modularity may, to some extent, reflect task modularity, it is still unclear whether data modularity directly corresponds with model modularity. One research avenue is to automate imposed data modularization regarding specific models in the spirit of AutoML. Similarly, automating task modularization can facilitate problem-solving and reduce human-introduced bias.

8 Conclusion

Deep learning is becoming dominant in many applications, such as computer vision and natural language processing. It is natural to ask ourselves whether there are guidelines for designing deep learning algorithms. Modularity is one guiding principle that has been put forward in the literature. This survey reveals that modularity is pervasive in three related yet distinct axes of deep learning: data, task, and model architecture. We observed that some modularity concepts come in the form of a prior, while others come in the form of a posterior.

The efforts of bringing the modularity principle into deep learning are not new; however, reviewing deep learning literature using the point of view of modularity is relatively new. This survey provides a step towards clarifying and investigating the notion of modularity in deep learning and elsewhere.

Acknowledgments

We gratefully acknowledge constructive feedback and suggestions from Birhanu Hailu Belay, Romain Egele, Felix Mohr, Hedi Tabia, and the reviewers. This work was supported by ChaLearn and the ANR (Agence Nationale de la Recherche, National Agency for Research) under AI chair of excellence HUMANIA, grant number ANR-19-CHIA-0022.

References

- [1] Accelerate Fast Math with Intel® oneAPI Math Kernel Library. <https://www.intel.com/content/www/us/en/developer/tools/oneapi>
- [2] Asmaa Abbas, Mohammed Abdelsamea, and Mohamed Gaber. DeTraC: Transfer Learning of Class Decomposed Medical Images in Convolutional Neural Networks. *IEEE Access*, PP:1–1, April 2020.
- [3] Madina Abdrakhmanova, Askat Kuzdeuov, Sheikh Jarju, Yerbolat Khassanov, Michael Lewis, and Huseyin Atakan Varol. Speakingfaces: A large-scale multimodal dataset of voice commands with visual and thermal video streams. *Sensors*, 21(10):3465, 2021.
- [4] Wickliffe C. Abraham and Anthony Robins. Memory retention – the synaptic stability versus plasticity dilemma. *Trends in Neurosciences*, 28(2):73–78, February 2005.
- [5] Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. *arXiv preprint arXiv:1610.01644*, 2016.
- [6] Ferran Alet, Tomás Lozano-Pérez, and Leslie P. Kaelbling. Modular meta-learning. *arXiv:1806.10166 [cs, stat]*, May 2019.
- [7] Anirudh Goyal ALIAS PARTH GOYAL, Aniket Didolkar, Nan Rosemary Ke, Charles Blundell, Philippe Beaudoin, Nicolas Heess, Michael C Mozer, and Yoshua Bengio. Neural production systems. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 25673–25687. Curran Associates, Inc., 2021.
- [8] Felipe Almeida and Geraldo Xexéo. Word Embeddings: A Survey, January 2019.
- [9] Mohammed Amer and Tomas Maul. A Review of Modularization Techniques in Artificial Neural Networks. *Artificial Intelligence Review*, 52, June 2019.
- [10] Ark Anderson, Kyle Shaffer, Artem Yankov, Court D. Corley, and Nathan O. Hodas. Beyond Fine Tuning: A Modular Approach to Learning on Small Data, November 2016.
- [11] Jacob Andreas. Measuring compositionality in representation learning. In *International Conference on Learning Representations*, 2019.
- [12] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural Module Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 39–48, Las Vegas, NV, USA, June 2016. IEEE.
- [13] G. Auda and M. Kamel. Modular neural networks a survey. *International journal of neural systems*, 9 2:129–51, 1999.
- [14] Jeremy Avigad. Modularity in mathematics. *The Review of Symbolic Logic*, 13(1):47–79, March 2020.
- [15] Farooq Azam. *Biologically Inspired Modular Neural Networks*. PhD thesis, Virginia Tech, May 2000.
- [16] Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. Systematic generalization: What is required and can it be learned? In *International Conference on Learning Representations*, 2019.
- [17] Carliss Y. Baldwin and Kim B. Clark. *Design Rules: The Power of Modularity*, volume 1. Cambridge, MA: MIT Press, first edition, 1999.
- [18] Randall Balestriero and Yann LeCun. POLICE: Provably Optimal Linear Constraint Enforcement for Deep Neural Networks, November 2022.
- [19] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018.
- [20] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT Pre-Training of Image Transformers. In *International Conference on Learning Representations*, 2022.
- [21] Paul Barham, Aakanksha Chowdhery, Jeff Dean, Sanjay Ghemawat, Steven Hand, Dan Hurt, Michael Isard, Hyeontaek Lim, Ruoming Pang, Sudip Roy, Brennan Saeta, Parker Schuh, Ryan Sepassi, Laurent El Shafey, Chandramohan A. Thekkath, and Yonghui Wu. Pathways: Asynchronous Distributed Dataflow for ML. *arXiv:2203.12533 [cs]*, March 2022.
- [22] H. Clark Barrett and Robert Kurzban. Modularity in cognition: Framing the debate. *Psychological Review*, 113(3):628–647, July 2006.
- [23] Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. *arXiv:1806.01261 [cs, stat]*, October 2018.

- [24] Erik J. Bekkers, Maxime W. Lafarge, Mitko Veta, Koen AJ Eppenhof, Josien PW Pluim, and Remco Duits. Roto-Translation Covariant Convolutional Networks for Medical Image Analysis. *arXiv:1804.03393 [cs, math]*, June 2018.
- [25] Birhanu Belay, Tewodros Habtegebrial, Marcus Liwicki, Gebeyehu Belay, and Didier Stricker. Factored Convolutional Neural Network for Amharic Character Image Recognition. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2906–2910, 2019.
- [26] Gabriel Béna and Dan F. M. Goodman. Extreme sparsity gives rise to functional specialization. *arXiv:2106.02626 [cs, q-bio]*, June 2021.
- [27] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional Computation in Neural Networks for faster models. *arXiv:1511.06297 [cs]*, January 2016.
- [28] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation. *arXiv:1308.3432 [cs]*, August 2013.
- [29] J. Bongard. Evolving modular genetic regulatory networks. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC’02 (Cat. No.02TH8600)*, volume 2, pages 1872–1877 vol.2, May 2002.
- [30] Diane Bouchacourt, Mark Ibrahim, and Ari Morcos. Grounding inductive biases in natural images: Invariance stems from variations in data. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 19566–19579. Curran Associates, Inc., 2021.
- [31] Nicholas Bourbaki. The Architecture of Mathematics. *The American Mathematical Monthly*, 57(4):221–232, April 1950.
- [32] Ulrik Brandes, Daniel Delling, Marco Gaertler, Robert Gorke, Martin Hoefer, Zoran Nikoloski, and Dorothea Wagner. On modularity clustering. *IEEE transactions on knowledge and data engineering*, 20(2):172–188, 2007.
- [33] Alexander Braylan, Mark Hollenbeck, Elliot Meyerson, and Risto Miikkulainen. Reuse of neural modules for general video game playing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [34] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature Verification using a “Siamese” Time Delay Neural Network. In J. Cowan, G. Tesauro, and J. Alspector, editors, *Advances in Neural Information Processing Systems*, volume 6. Morgan-Kaufmann, 1994.
- [35] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [36] Qingqing Cao, Harsh Trivedi, Aruna Balasubramanian, and Niranjan Balasubramanian. DeFormer: Decomposing Pre-trained Transformers for Faster Question Answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4497, Online, July 2020. Association for Computational Linguistics.
- [37] Richard G Casey and Eric Lecolinet. A survey of methods and strategies in character segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 18(7):690–706, 1996.
- [38] Michael Chang, Abhishek Gupta, Sergey Levine, and Thomas L. Griffiths. Automatically composing representation transformations as a means for generalization. In *International Conference on Learning Representations*, 2019.
- [39] Michael Chang, Sid Kaushik, Sergey Levine, and Tom Griffiths. Modularity in Reinforcement Learning via Algorithmic Independence in Credit Assignment. In *International Conference on Machine Learning*, pages 1452–1462. PMLR, July 2021.
- [40] Shuxiao Chen, Edgar Dobriban, and Jane Lee. A group-theoretic framework for data augmentation. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21321–21333. Curran Associates, Inc., 2020.
- [41] Xiaoxue Chen, Lianwen Jin, Yuanzhi Zhu, Canjie Luo, and Tianwei Wang. Text Recognition in the Wild: A Survey. *arXiv:2005.03492 [cs]*, December 2020.
- [42] Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15745–15753, Nashville, TN, USA, June 2021. IEEE.
- [43] Yutian Chen, Abram L Friesen, Feryal Behbahani, Arnaud Doucet, David Budden, Matthew Hoffman, and Nando de Freitas. Modular meta-learning with shrinkage. *Advances in Neural Information Processing Systems*, 33:2858–2869, 2020.

- [44] Guillaume Chevalier. Long short-term memory (LSTM cell). *Wikipedia*, September 2022.
- [45] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–Decoder approaches. *Syntax, Semantics and Structure in Statistical Translation*, page 103, 2014.
- [46] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1251–1258, 2017.
- [47] Noam Chomsky. *Aspects of the Theory of Syntax*. MIT Press, 1965.
- [48] Amit Choudhary, Rahul Rishi, and Savita Ahlawat. A new character segmentation approach for off-line cursive handwritten words. *Procedia Computer Science*, 17:88–95, 2013.
- [49] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. PaLM: Scaling Language Modeling with Pathways. *arXiv:2204.02311 [cs]*, April 2022.
- [50] Brian Chu, Vashisht Madhavan, Oscar Beijbom, Judy Hoffman, and Trevor Darrell. Best practices for fine-tuning visual classifiers to new domains. In *European Conference on Computer Vision*, pages 435–442. Springer, 2016.
- [51] Ignasi Clavera, David Held, and Pieter Abbeel. Policy transfer via modularity and reward guiding. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1537–1544. IEEE, 2017.
- [52] Jeff Clune, Jean-Baptiste Mouret, and Hod Lipson. The evolutionary origins of modularity. *Proceedings of the Royal Society b: Biological sciences*, 280(1755):20122863, 2013.
- [53] Taco S. Cohen and Max Welling. Group Equivariant Convolutional Networks. *arXiv:1602.07576 [cs, stat]*, June 2016.
- [54] Taco S. Cohen and Max Welling. Steerable CNNs. *arXiv:1612.08498 [cs, stat]*, December 2016.
- [55] Sarah Cohen-Boulakia, Khalid Belhajjame, Olivier Collin, Jérôme Chopard, Christine Froidevaux, Alban Gaignard, Konrad Hinsén, Pierre Larmande, Yvan Le Bras, Frédéric Lemoine, et al. Scientific workflows for computational reproducibility in the life sciences: Status, challenges and opportunities. *Future Generation Computer Systems*, 75:284–298, 2017.
- [56] Cédric Colas, Pierre Fournier, Mohamed Chetouani, Olivier Sigaud, and Pierre-Yves Oudeyer. Curious: Intrinsically motivated modular multi-goal reinforcement learning. In *International Conference on Machine Learning*, pages 1331–1340. PMLR, 2019.
- [57] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to Algorithms*. 2022.
- [58] Leda Cosmides and John Tooby. Cognitive Adaptations for Social Exchange. *undefined*, pages 163–228, 1992.
- [59] Leda Cosmides and John Tooby. Origins of domain specificity: The evolution of functional organization. In Lawrence A. Hirschfeld and Susan A. Gelman, editors, *Mapping the Mind: Domain Specificity in Cognition and Culture*, pages 85–116. Cambridge University Press, Cambridge, 1994.
- [60] Róbert Csordás, Kazuki Irie, and Jürgen Schmidhuber. CTL++: Evaluating Generalization on Never-Seen Compositional Patterns of Known Functions, and Compatibility of Neural Representations. In *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, December 2022.
- [61] Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Are Neural Nets Modular? Inspecting Functional Modularity Through Differentiable Weight Masks. In *International Conference on Learning Representations*, 2021.
- [62] Vanessa D’Amario, Tomotake Sasaki, and Xavier Boix. How Modular should Neural Module Networks Be for Systematic Generalization? In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [63] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [64] Coline Devin, Abhishek Gupta, Trevor Darrell, Pieter Abbeel, and Sergey Levine. Learning modular neural network policies for multi-task and multi-robot transfer. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2169–2176. IEEE, 2017.
- [65] Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu. Exploiting Cyclic Symmetry in Convolutional Neural Networks. *arXiv:1602.02660 [cs]*, May 2016.

- [66] Changxing Ding and Dacheng Tao. Trunk-branch ensemble convolutional neural networks for video-based face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):1002–1014, 2017.
- [67] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2021.
- [68] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR, 2022.
- [69] Cian Eastwood and Christopher K. I. Williams. A Framework for the Quantitative Evaluation of Disentangled Representations. In *Sixth International Conference on Learning Representations (ICLR 2018)*, May 2018.
- [70] David Eigen, Marc’Aurelio Ranzato, and Ilya Sutskever. Learning factored representations in a deep mixture of experts. In *ICLR Workshop*, 2014.
- [71] Adrian El Baz, Ihsan Ullah, Edesio Alcobaça, André C. P. L. F. Carvalho, Hong Chen, Fabio Ferreira, Henry Gouk, Chaoyu Guan, Isabelle Guyon, Timothy Hospedales, Shell Hu, Mike Huisman, Frank Hutter, Zhengying Liu, Felix Mohr, Ekrem Öztürk, Jan N. van Rijn, Haozhe Sun, Xin Wang, and Wenwu Zhu. Lessons learned from the NeurIPS 2021 MetaDL challenge: Backbone fine-tuning without episodic meta-learning dominates for few-shot learning image classification. In Douwe Kiela, Marco Ciccone, and Barbara Caputo, editors, *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pages 80–96. PMLR, December 2022.
- [72] Kai Olav Ellefsen, Jean-Baptiste Mouret, and Jeff Clune. Neural Modularity Helps Organisms Evolve to Learn New Skills without Forgetting Old Skills. *PLOS Computational Biology*, 11(4):e1004128, April 2015.
- [73] Gamaleldin F. Elsayed, Prajit Ramachandran, Jonathon Shlens, and Simon Kornblith. Revisiting Spatial Invariance with Low-Rank Local Connectivity. *arXiv:2002.02959 [cs, stat]*, August 2020.
- [74] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural Architecture Search. pages 69–86.
- [75] William Fedus, Jeff Dean, and Barret Zoph. A Review of Sparse Expert Models in Deep Learning, September 2022.
- [76] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- [77] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. PathNet: Evolution Channels Gradient Descent in Super Neural Networks. *arXiv:1701.08734 [cs]*, January 2017.
- [78] Daniel Filan, Stephen Casper, Shlomi Hod, Cody Wild, Andrew Critch, and Stuart Russell. Clusterability in Neural Networks. *arXiv:2103.03386 [cs]*, March 2021.
- [79] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv:1703.03400 [cs]*, July 2017.
- [80] Jerry A. Fodor. The Modularity of Mind. April 1983.
- [81] Jerry A. Fodor. *The Mind Doesn’t Work That Way: The Scope and Limits of Computational Psychology*. MIT Press, 2000.
- [82] Jerry A Fodor and Zenon W Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988.
- [83] Martin Ford. *Architects of Intelligence: The Truth about AI from the People Building It*. Packt Publishing, Birmingham, UK, first published: november 2018 edition, 2018.
- [84] Willem E. Frankenhuis and Annemie Ploeger. Evolutionary Psychology Versus Fodor: Arguments For and Against the Massive Modularity Hypothesis. *Philosophical Psychology*, 20(6):687–710, December 2007.
- [85] Robert French. Using Semi-Distributed Representations to Overcome Catastrophic Forgetting in Connectionist Networks. 1991.
- [86] Ninnart Fuengfusin and Hakaru Tamukoh. Network with Sub-networks: Layer-wise Detachable Neural Network. *Journal of Robotics, Networking and Artificial Life*, 7(4):240–244, 2020.
- [87] Tomer Galanti and Lior Wolf. On the Modularity of Hypernetworks. *arXiv:2002.10006 [cs, stat]*, November 2020.
- [88] Hongyang Gao and Shuiwang Ji. Efficient and Invariant Convolutional Neural Networks for Dense Prediction. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 871–876, 2017.

- [89] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image Style Transfer Using Convolutional Neural Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, Las Vegas, NV, USA, June 2016. IEEE.
- [90] Pralhad Gavali and J. Saira Banu. Chapter 6 - deep convolutional neural network for image classification on CUDA platform. In Arun Kumar Sangaiah, editor, *Deep Learning and Parallel Computing Environment for Bioengineering Systems*, pages 99–122. Academic Press, 2019.
- [91] Peter Gentile. Theory of Modularity, a Hypothesis. *Procedia Computer Science*, 20, December 2013.
- [92] Badih Ghazi, Rina Panigrahy, and Joshua Wang. Recursive Sketches for Modular Deep Learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2211–2220. PMLR, May 2019.
- [93] Daniel Gómez, J. Tinguaro Rodríguez, Javier Yáñez, and Javier Montero. A new modularity measure for Fuzzy Community detection problems based on overlap and grouping functions. *International Journal of Approximate Reasoning*, 74:88–107, July 2016.
- [94] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [95] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]*, June 2014.
- [96] Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. In *International Conference on Learning Representations*, 2021.
- [97] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913, 2017.
- [98] Scott Gray, Alec Radford, and Diederik P Kingma. GPU Kernels for Block-Sparse Weights. Technical report.
- [99] David Ha, Andrew Dai, and Quoc V. Le. HyperNetworks. *arXiv:1609.09106 [cs]*, December 2016.
- [100] Guy Hachohen and Daphna Weinshall. On The Power of Curriculum Learning in Training Deep Networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2535–2544. PMLR, June 2019.
- [101] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, second edition, 2009.
- [102] Jiaao He, Jidong Zhai, Tiago Antunes, Haojie Wang, Fuwen Luo, Shangfeng Shi, and Qin Li. FasterMoE: Modeling and optimizing training of large-scale dynamic pre-trained models. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 120–134, 2022.
- [103] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [104] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4918–4927, 2019.
- [105] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [106] G. E. Hinton and R. R. Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786):504–507, 2006.
- [107] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [108] Michel A Hofman. Evolution of the human brain: When bigger is better. *Frontiers in neuroanatomy*, 8:15, 2014.
- [109] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [110] Guosheng Hu, Yuxin Hu, Kai Yang, Zehao Yu, Flood Sung, Zhihong Zhang, Fei Xie, Jianguo Liu, Neil Robertson, Timothy Hospedales, and Qiangwei Miemie. Deep Stock Representation Learning: From Candlestick Charts to Investment Decisions. *arXiv:1709.03803 [q-fin]*, February 2018.
- [111] Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to Reason: End-to-End Module Networks for Visual Question Answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.

- [112] Jing Huang, Guan Pang, Rama Kovvuri, Mandy Toh, Kevin J Liang, Praveen Krishnan, Xi Yin, and Tal Hassner. A multiplexed network for end-to-end, multilingual OCR. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4547–4557, 2021.
- [113] Joost Huizinga, Jeff Clune, and Jean-Baptiste Mouret. Evolving neural networks that are both modular and regular: Hyperneat plus the connection cost technique. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 697–704, 2014.
- [114] Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795, 2020.
- [115] Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, Dennis Ulmer, Florian Schottmann, Khuyagbaatar Batsuren, Kaiser Sun, Koustuv Sinha, Leila Khalatbari, Maria Ryskina, Rita Frieske, Ryan Cotterell, and Zhijing Jin. State-of-the-art generalisation research in NLP: A taxonomy and review, October 2022.
- [116] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren, editors. *Automatic Machine Learning: Methods, Systems, Challenges*. Springer, 2019.
- [117] Riashat Islam, Hongyu Zang, Anirudh Goyal, Alex Lamb, Kenji Kawaguchi, Xin Li, Romain Laroche, Yoshua Bengio, and Remi Tachet Des Combes. Discrete Factorial Representations as an Abstraction for Goal Conditioned Reinforcement Learning, October 2022.
- [118] Robert A Jacobs, Michael I Jordan, and Andrew G Barto. Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive science*, 15(2):219–250, 1991.
- [119] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1):79–87, March 1991.
- [120] Khurram Javed and Martha White. Meta-Learning Representations for Continual Learning. *arXiv:1905.12588 [cs, stat]*, October 2019.
- [121] Tian Jin and Seokin Hong. Split-CNN: Splitting Window-Based Operations in Convolutional Neural Networks for Memory System Optimization. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS ’19, pages 835–847, New York, NY, USA, 2019. Association for Computing Machinery.
- [122] Li Jing, Jiachen Zhu, and Yann LeCun. Masked Siamese ConvNets, June 2022.
- [123] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6(2):181–214, 1994.
- [124] Cheng Ju, Aurélien Bibaut, and Mark van der Laan. The relative performance of ensemble methods with deep convolutional neural networks for image classification. *Journal of Applied Statistics*, 45(15):2800–2818, 2018.
- [125] Dan Jurafsky and James H Martin. *Speech and language processing* (3rd draft ed.), 2019.
- [126] Menelaos Kanakis, David Bruggemann, Suman Saha, Stamatios Georgoulis, Anton Obukhov, and Luc Van Gool. Reparameterizing convolutions for incremental multi-task learning without task interference. In *European Conference on Computer Vision*, pages 689–707. Springer, 2020.
- [127] Nora Kassner, Oyvind Tafjord, Hinrich Schütze, and Peter Clark. BeliefBank: Adding Memory to a Pre-Trained Language Model for a Systematic Notion of Belief. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8849–8861, 2021.
- [128] Amandeep Kaur, Seema Baghla, and Sunil Kumar. Study of various character segmentation techniques for handwritten off-line cursive words: A review. *International Journal of Advances in Science Engineering and Technology*, 3(3):154–158, 2015.
- [129] Zixuan Ke, Bing Liu, Nianzu Ma, Hu Xu, and Lei Shu. Achieving forgetting prevention and knowledge transfer in continual learning. *Advances in Neural Information Processing Systems*, 34:22443–22456, 2021.
- [130] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [131] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. In *ICLR*, 2017.
- [132] Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. Measuring compositional generalization: A comprehensive method on realistic data. In *International Conference on Learning Representations*, 2020.

- [133] Juyong Kim, Yookoon Park, Gunhee Kim, and Sung Ju Hwang. SplitNet: Learning to Semantically Split Deep Networks for Parameter Reduction and Model Parallelization. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1866–1874. PMLR, July 2017.
- [134] Hiroaki Kingetsu, Kenichi Kobayashi, and Taiji Suzuki. Neural Network Module Decomposition and Recomposition, December 2021.
- [135] Louis Kirsch, Julius Kunze, and David Barber. Modular Networks: Learning to Decompose Neural Computation. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [136] Eunjeong Koh and Shlomo Dubnov. Comparison and Analysis of Deep Audio Embeddings for Music Emotion Recognition, April 2021.
- [137] Yamuna Krishnamurthy and Chris Watkins. Interpretability in gated modular neural networks. In *eXplainable AI Approaches for Debugging and Diagnosis.*, 2021.
- [138] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, volume 25, 2012.
- [139] David Krueger, Tegan Maharaj, Janos Kramar, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke, Anirudh Goyal, Yoshua Bengio, Aaron Courville, and Christopher Pal. Zoneout: Regularizing RNNs by randomly preserving hidden activations. In *International Conference on Learning Representations*, 2017.
- [140] Ray Kurzweil. *How to Create a Mind: The Secret of Human Thought Revealed*. Penguin Books, USA, 2013.
- [141] Steinar Laenen and Luca Bertinetto. On episodes, prototypical networks, and few-shot learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 24581–24592. Curran Associates, Inc., 2021.
- [142] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, December 2015.
- [143] Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, pages 2873–2882. PMLR, 2018.
- [144] Brenden M. Lake. Compositional generalization through meta sequence-to-sequence learning. *arXiv:1906.05381 [cs]*, October 2019.
- [145] Y. LeCun, Fu Jie Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pages II–104 Vol.2, June 2004.
- [146] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. 2022.
- [147] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [148] Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. *Advances in neural information processing systems*, 2, 1989.
- [149] Ming Li and Paul M.B. Vitnyi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer Publishing Company, Incorporated, third edition, 2008.
- [150] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural speech synthesis with transformer network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6706–6713, 2019.
- [151] Zhi Li, Bo Wu, Qi Liu, Likang Wu, Hongke Zhao, and Tao Mei. Learning the compositional visual coherence for complementary recommendations. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3536–3543. International Joint Conferences on Artificial Intelligence Organization, July 2020.
- [152] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *International Conference on Learning Representations*, 2018.
- [153] João Loula, Marco Baroni, and Brenden M. Lake. Rearranging the familiar: Testing compositional generalization in recurrent networks. In *BlackboxNLP@EMNLP*, pages 108–114, 2018.
- [154] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. Disentangled graph convolutional networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4212–4221. PMLR, June 2019.
- [155] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive Single-Tasking of Multiple Tasks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1851–1860, Long Beach, CA, USA, June 2019. IEEE.

- [156] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015.
- [157] Nicolas Y. Masse, Gregory D. Grant, and David J. Freedman. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*, 115(44):E10467–E10475, 2018.
- [158] Vittorio Mazzia, Francesco Salvetti, and Marcello Chiaberge. Efficient-capsnet: Capsule network with self-attention routing. *Scientific reports*, 11(1):1–13, 2021.
- [159] David McNeely-White, J. Ross Beveridge, and Bruce A. Draper. Inception and ResNet features are (almost) equivalent. *Cognitive Systems Research*, 59:312–318, January 2020.
- [160] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and Editing Factual Associations in GPT. February 2022.
- [161] Elliot Meyerson and Risto Miikkulainen. Modular universal reparameterization: Deep multi-task learning across diverse domains. *Advances in Neural Information Processing Systems*, 32, 2019.
- [162] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. Fast Model Editing at Scale. *arXiv:2110.11309 [cs]*, October 2021.
- [163] Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. Memory-Based Model Editing at Scale. In *International Conference on Machine Learning*, 2022.
- [164] Sarthak Mittal, Yoshua Bengio, and Guillaume Lajoie. Is a Modular Architecture Enough?, 2022.
- [165] Sarthak Mittal, Sharath Chandra Raparthy, Irina Rish, Yoshua Bengio, and Guillaume Lajoie. Compositional attention: Disentangling search and retrieval. In *International Conference on Learning Representations*, 2022.
- [166] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous Methods for Deep Reinforcement Learning. *arXiv e-prints*, page arXiv:1602.01783, February 2016.
- [167] Vladimir Modrak and Zuzana Soltysova. Development of the Modularity Measure for Assembly Process Structures. *Mathematical Problems in Engineering*, 2021:e4900748, December 2021.
- [168] Stefanie Muff, Francesco Rao, and Amedeo Caglisch. Local modularity measure for network clusterizations. *Physical Review E*, 72(5):056107, November 2005.
- [169] Shikhar Murty, Pratyusha Sharma, Jacob Andreas, and Christopher D. Manning. Characterizing Intrinsic Compositionality in Transformers with Tree Projections, November 2022.
- [170] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
- [171] Michael Opitz, Horst Possegger, and Horst Bischof. Efficient model averaging for deep neural networks. In *Asian Conference on Computer Vision*, pages 205–220. Springer, 2016.
- [172] Oleksiy Ostapenko, Pau Rodriguez, Massimo Caccia, and Laurent Charlin. Continual learning via local module composition. *Advances in Neural Information Processing Systems*, 34:30298–30312, 2021.
- [173] Oleksiy Ostapenko, Pau Rodriguez, Alexandre Lacoste, and Laurent Charlin. Attention for compositional modularity. In *NeurIPS ’22 Workshop on All Things Attention: Bridging Different Perspectives on Attention*, 2022.
- [174] Rangeet Pan and Hridesh Rajan. On Decomposing a Deep Neural Network into Modules. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2020, pages 889–900, New York, NY, USA, 2020. Association for Computing Machinery.
- [175] Rangeet Pan and Hridesh Rajan. Decomposing Convolutional Neural Networks into Reusable and Replaceable Modules. In *Proceedings of The 44th International Conference on Software Engineering (ICSE 2022)*, December 2021.
- [176] Giambattista Parascandolo, Niki Kilbertus, Mateo Rojas-Carulla, and Bernhard Schölkopf. Learning independent causal mechanisms. In *International Conference on Machine Learning*, pages 4036–4044. PMLR, 2018.
- [177] D. L. Parnas. On the criteria to be used in decomposing systems into modules. *Communications of the ACM*, 15(12):1053–1058, December 1972.

- [178] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [179] Deepak Pathak, Christopher Lu, Trevor Darrell, Phillip Isola, and Alexei A Efros. Learning to control self-assembling morphologies: A study of generalization via modularity. *Advances in Neural Information Processing Systems*, 32, 2019.
- [180] Jose B Pereira-Leal, Emmanuel D Levy, and Sarah A Teichmann. The origins and evolution of functional modules: Lessons from protein complexes. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1467):507–517, 2006.
- [181] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. Adaptive Computation and Machine Learning Series. MIT Press, Cambridge, MA, USA, November 2017.
- [182] Timothée Poisot. An a posteriori measure of network modularity. *F1000Research*, 2:130, December 2013.
- [183] Edoardo Ponti. *Inductive Bias and Modular Design for Sample-Efficient Neural Language Learning*. PhD thesis, University of Cambridge, 2021.
- [184] Edoardo M. Ponti, Alessandro Sordoni, Yoshua Bengio, and Siva Reddy. Combining Modular Skills in Multitask Learning, March 2022.
- [185] Senthil Purushwalkam, Maximilian Nickel, Abhinav Gupta, and Marc’Aurelio Ranzato. Task-driven modular networks for zero-shot compositional learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3593–3602, 2019.
- [186] Zenon Pylyshyn. Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3):341–365, June 1999.
- [187] Jun-Fei Qiao, Xi Meng, Wen-Jing Li, and Bogdan M. Wilamowski. A novel modular RBF neural network based on a brain-like partition method. *Neural Computing and Applications*, 32(3):899–911, February 2020.
- [188] Nasim Rahaman, Muhammad Waleed Gondal, Shruti Joshi, Peter Gehler, Yoshua Bengio, Francesco Locatello, and Bernhard Schölkopf. Dynamic inference with neural interpreters. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 10985–10998. Curran Associates, Inc., 2021.
- [189] Prajit Ramachandran and Quoc V. Le. Diversity and depth in per-example routing models. In *International Conference on Learning Representations*, 2019.
- [190] G Ranganathan et al. A study to find facts behind preprocessing on deep learning algorithms. *Journal of Innovative Image Processing (JIIP)*, 3(01):66–74, 2021.
- [191] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- [192] J. Reisinger, K. Stanley, and R. Miikkulainen. Evolving Reusable Neural Modules. In *GECCO*, 2004.
- [193] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021.
- [194] Karl Ridgeway and Michael C Mozer. Learning Deep Disentangled Embeddings With the F-Statistic Loss. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [195] Philip Robbins. Modularity of Mind. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 edition, 2017.
- [196] John S Rose. *A Course on Group Theory*. Courier Corporation, 1994.
- [197] Clemens Rosenbaum, Ignacio Cases, Matthew Riemer, and Tim Klinger. Routing Networks and the Challenges of Modular and Compositional Computation, April 2019.
- [198] Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing Networks: Adaptive Selection of Non-Linear Functions for Multi-Task Learning. In *International Conference on Learning Representations*, 2018.
- [199] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- [200] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.

- [201] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet Large Scale Visual Recognition Challenge. In *International Journal of Computer Vision*, volume 115, pages 211–252. Springer, 2015.
- [202] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive Neural Networks. *arXiv:1606.04671 [cs]*, September 2016.
- [203] Guillaume Salha-Galvan, Johannes F. Lutzeyer, George Dasoulas, Romain Hennequin, and Michalis Vazirgiannis. Modularity-Aware Graph Autoencoders for Joint Community Detection and Link Prediction, June 2022.
- [204] M. Schenkel, H. Weissman, I. Guyon, C. Nohl, and D. Henderson. Recognition-based segmentation of on-line hand-printed words. In S. Hanson, J. Cowan, and C. Giles, editors, *Advances in Neural Information Processing Systems*, volume 5. Morgan-Kaufmann, 1992.
- [205] Melissa Schilling. Toward a General Modular Systems Theory and Its Application to Interfirm Product Modularity. *Academy of Management Review*, 25, April 2000.
- [206] Jürgen Schmidhuber. Towards compositional learning in dynamic networks. 1990.
- [207] Albrecht Schmidt and Zuhair Bandar. Modularity - A Concept For New Neural Network Architectures. November 2001.
- [208] Yue Shao and Victor M. Zavala. Modularity measures: Concepts, computation, and applications to manufacturing systems. *AIChE Journal*, 66(6):e16965, 2020.
- [209] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [210] Baoguang Shi, Xiang Bai, and Cong Yao. Script identification in the wild via discriminative convolutional neural network. *Pattern Recognition*, 52:448–458, April 2016.
- [211] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual Learning with Deep Generative Replay. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [212] Hiroaki Shiokawa, Yasuhiro Fujiwara, and Makoto Onizuka. Fast algorithm for modularity-based graph clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, pages 1170–1176, 2013.
- [213] L Sifre. *Rigid-Motion Scattering for Image Classification [PhD Thesis]*. PhD thesis, 2014.
- [214] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016.
- [215] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, October 2017.
- [216] P. Y. Simard, D. Steinkraus, and J. C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, pages 958–963, August 2003.
- [217] Herbert A. Simon. The Architecture of Complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.
- [218] Herbert A. Simon and Albert Ando. Aggregation of variables in dynamic systems. *Econometrica*, 29(2):111–138, 1961.
- [219] Christopher Simpkins and Charles Isbell. Composable modular reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4975–4982, 2019.
- [220] Anton Sinitsin, Vsevolod Plohotnyuk, Dmitry Pyrkin, Sergei Popov, and Artem Babenko. Editable Neural Networks. In *International Conference on Learning Representations*, 2019.
- [221] Samuel L Smith, Pieter-Jan Kindermans, Chris Ying, and Quoc V Le. Don’t decay the learning rate, increase the batch size. In *International Conference on Learning Representations*, 2018.

- [222] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhunoye, George Zerveas, Vijay Korthikanti, Elton Zhang, Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi, Yuxiong He, Michael Houston, Saurabh Tiwary, and Bryan Catanzaro. Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model, February 2022.
- [223] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical Networks for Few-shot Learning. *arXiv:1703.05175 [cs, stat]*, June 2017.
- [224] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- [225] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 843–852, 2017.
- [226] Guolei Sun, Thomas Probst, Danda Pani Paudel, Nikola Popovic, Menelaos Kanakis, Jagruti Patel, Dengxin Dai, and Luc Van Gool. Task Switching Network for Multi-task Learning. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8271–8280, Montreal, QC, Canada, October 2021. IEEE.
- [227] Haozhe Sun, Wei-Wei Tu, and Isabelle M. Guyon. OmniPrint: A Configurable Printed Character Synthesizer. In *Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [228] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- [229] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [230] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [231] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv:1703.01780 [cs, stat]*, April 2018.
- [232] Surat Teerapittayanon, Bradley McDanel, and H.T. Kung. BranchyNet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469, 2016.
- [233] Alexander Terekhov, Guglielmo Montone, and J. O’Regan. Knowledge Transfer in Deep Block-Modular Neural Networks. In *Biomimetic and Biohybrid Systems*, pages 268–279. Springer, July 2015.
- [234] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5, 2015.
- [235] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *International Conference on Learning Representations*, 2019.
- [236] Ihsan Ullah, Dustin Carrion, Sergio Escalera, Isabelle M. Guyon, Mike Huisman, Felix Mohr, Jan N. van Rijn, Haozhe Sun, Joaquin Vanschoren, and Phan Anh Vu. Meta-Album: Multi-domain Meta-Dataset for Few-Shot Image Classification, 2022.
- [237] Ivan I Vankov and Jeffrey S Bowers. Training neural networks to encode symbols enables combinatorial generalization. *Philosophical Transactions of the Royal Society B*, 375(1791):20190309, 2020.
- [238] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [239] Tom Veniat, Ludovic Denoyer, and Marc’Aurelio Ranzato. Efficient Continual Learning with Modular Networks and Task-Driven Priors. In *9th International Conference on Learning Representations, ICLR 2021*, 2021.
- [240] Ulrike von Luxburg, Robert C. Williamson, and Isabelle Guyon. Clustering: Science or Art? In *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, pages 65–79. JMLR Workshop and Conference Proceedings, June 2012.
- [241] Gunter P. Wagner and Lee Altenberg. Perspective: Complex Adaptations and the Evolution of Evolvability. *Evolution*, 50(3):967–976, 1996.
- [242] Haoxiang Wang, Han Zhao, and Bo Li. Bridging multi-task learning and meta-learning: Towards efficient training and effective adaptation. In *International Conference on Machine Learning*, pages 10991–11002. PMLR, 2021.

- [243] Jianan Wang, Eren Sezener, David Budden, Marcus Hutter, and Joel Veness. A Combinatorial Perspective on Transfer Learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 918–929. Curran Associates, Inc., 2020.
- [244] Ruohan Wang, Massimiliano Pontil, and Carlo Ciliberto. The role of global labels in few-shot classification and how to infer them. In *Advances in Neural Information Processing Systems*, volume 34, pages 27160–27170, 2021.
- [245] Chihiro Watanabe, Kaoru Hiramatsu, and Kunio Kashino. Modular representation of layered neural networks. *Neural Networks*, 97:62–73, 2018.
- [246] Maurice Weiler and Gabriele Cesa. General $S(2)$ -Equivariant Steerable CNNs. *arXiv:1911.08251 [cs, eess]*, April 2021.
- [247] Maurice Weiler, Fred A. Hamprecht, and Martin Storath. Learning Steerable Filters for Rotation Equivariant CNNs. *arXiv:1711.07289 [cs]*, March 2018.
- [248] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow. Harmonic Networks: Deep Translation and Rotation Equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5028–5037, 2017.
- [249] Likang Wu, Zhi Li, Hongke Zhao, Qi Liu, Jun Wang, Mengdi Zhang, and Enhong Chen. Learning the implicit semantic representation on graph-structured data. In *International Conference on Database Systems for Advanced Applications*, pages 3–19. Springer, 2021.
- [250] Yuhuai Wu, Elman Mansimov, Shun Liao, Roger Grosse, and Jimmy Ba. Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. *arXiv e-prints*, page arXiv:1708.05144, August 2017.
- [251] Saining Xie, Ross Girshick, Piotr Dollar, Zhuowen Tu, and Kaiming He. Aggregated Residual Transformations for Deep Neural Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [252] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring Randomly Wired Neural Networks for Image Recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [253] Chao Xiong, Xiaowei Zhao, Danhang Tang, Karlekar Jayashree, Shuicheng Yan, and Tae-Kyun Kim. Conditional convolutional neural network for modality-aware face recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3667–3675, 2015.
- [254] I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *CoRR*, abs/1905.00546, 2019.
- [255] Shudong Yang, Xueying Yu, and Ying Zhou. LSTM and GRU neural network performance comparison study: Taking yelp review dataset as an example. In *2020 International Workshop on Electronic Communication and Artificial Intelligence (IWECAI)*, pages 98–101, 2020.
- [256] Bangpeng Yao, Dirk Walther, Diane Beck, and Li Fei-fei. Hierarchical mixture of classification experts uncovers interactions between brain regions. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009.
- [257] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, pages 7105–7114. PMLR, 2019.
- [258] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *arXiv:1411.1792 [cs]*, November 2014.
- [259] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable Neural Networks. In *International Conference on Learning Representations*, 2019.
- [260] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MAttNet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315, 2018.
- [261] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient Surgery for Multi-Task Learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 5824–5836. Curran Associates, Inc., 2020.
- [262] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [263] Julian Zaidi, Jonathan Boilard, Ghyslain Gagnon, and Marc-André Carbonneau. Measuring Disentanglement: A Review of Metrics. *arXiv:2012.09276 [cs]*, January 2021.

- [264] Quanshi Zhang, Yu Yang, Qian Yu, and Ying Nian Wu. Network Transplanting. *arXiv:1804.10272 [cs, stat]*, December 2018.
- [265] Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [266] Allan Zhou, Tom Knowles, and Chelsea Finn. Meta-Learning Symmetries by Reparameterization. *arXiv:2007.02933 [cs, stat]*, October 2020.
- [267] Tianyi Zhou, Shengjie Wang, and Jeff A Bilmes. Diverse ensemble evolution: Curriculum data-model marriage. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [268] Zhi-Hua Zhou. *Ensemble Methods: Foundations and Algorithms*. CRC press, 2012.
- [269] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Computer Vision (ICCV), 2017 IEEE International Conference On*, 2017.
- [270] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning Transferable Architectures for Scalable Image Recognition. *arXiv:1707.07012 [cs, stat]*, April 2018.