
A SURVEY ON KNOWLEDGE EDITING OF NEURAL NETWORKS

Vittorio Mazzia
Alexa AI, Amazon
vmazzia@amazon.com

Alessandro Pedrani
Alexa AI, Amazon
pedrana@amazon.com

Andrea Caciolai
Alexa AI, Amazon
andccl@amazon.com

Kay Rottmann
Alexa AI, Amazon
krrothm@amazon.com

Davide Bernardi
Alexa AI, Amazon
dvdbe@amazon.com

ABSTRACT

Deep neural networks are becoming increasingly pervasive in academia and industry, matching and surpassing human performance on a wide variety of fields and related tasks. However, just as humans, even the largest artificial neural networks make mistakes, and once-correct predictions can become invalid as the world progresses in time. Augmenting datasets with samples that account for mistakes or up-to-date information has become a common workaround in practical applications. However, the well-known phenomenon of catastrophic forgetting poses a challenge in achieving precise changes in the implicitly memorized knowledge of neural network parameters, often requiring a full model re-training to achieve desired behaviors. That is expensive, unreliable, and incompatible with the current trend of large self-supervised pre-training, making it necessary to find more efficient and effective methods for adapting neural network models to changing data. To address this need, knowledge editing is emerging as a novel area of research that aims to enable reliable, data-efficient, and fast changes to a pre-trained target model, without affecting model behaviors on previously learned tasks. In this survey, we provide a brief review of this recent artificial intelligence field of research. We first introduce the problem of editing neural networks, formalize it in a common framework and differentiate it from more notorious branches of research such as continuous learning. Next, we provide a review of the most relevant knowledge editing approaches and datasets proposed so far, grouping works under four different families: regularization techniques, meta-learning, direct model editing, and architectural strategies. Finally, we outline some intersections with other fields of research and potential directions for future works.

Keywords Knowledge Editing · Model Editing · Neural Networks Editing · Continual Learning

1 Introduction

In stark contrast to artificial neural networks (ANN), [Cichon and Gan \(2015\)](#), humans and other animals seem capable of learning and editing their knowledge continuously. Indeed, literature studies indicate that the mammalian brain could prevent catastrophic forgetting [Ratcliff \(1990\)](#) by safeguarding previously acquired knowledge, thereby reducing the plasticity of a proportion of synapses and ensuring their long-term stability [Benna and Fusi \(2016\)](#); [Yang et al. \(2009\)](#); [Cichon and Gan \(2015\)](#). On the contrary, ANNs not only struggle to learn new tasks in a sequential fashion [Kirkpatrick et al. \(2017\)](#), but also edit acquired knowledge on the same data distribution and task [Huang et al. \(2023\)](#). Indeed, unlike conventional knowledge base systems that explicitly store knowledge, neural models implicitly memorize facts and tasks in their parameters, making it difficult to directly access and interpret their computation and memories [Voita et al. \(2019\)](#); [Belinkov and Glass \(2019\)](#). Making even minor modifications can lead to a decrease in performance on previously learnt tasks, or even cause the entire computation to fail due to the well-documented issue of catastrophic forgetting [Ratcliff \(1990\)](#). Therefore, modifying their acquired knowledge is a challenging problem.

Just as humans, ANNs make mistakes and as we trust them with increasingly important decisions, the cost of such mistakes grows ever higher [Sinitsin et al. \(2020\)](#). Therefore, given that mistakes are inevitable, it is crucial for deep

learning practitioners to possess the ability to adjust model behaviors by promptly correcting errors as they surface. Currently, practical applications employing deep learning techniques have been relying on different workarounds to tackle this problem. In particular, a full re-training using datasets augmented with samples that account for the mistakes or up-to-date information is a common choice [Sinitstin et al. \(2020\)](#). The endeavor needed for fine-tuning atop pre-trained models [Sarzynska-Wawer et al. \(2021\)](#); [Devlin et al. \(2018\)](#); [Oquab et al. \(2023\)](#); [Weiss et al. \(2016\)](#) is frequently substantiated by the diminished dataset size and computational resources needed. On the other hand, this is not always true and manually curated, deterministic mechanisms that overrule model predictions on problematic samples can be the preferred choice [Sinitstin et al. \(2020\)](#). However, while being simple, this second approach is fully localized and not robust to factor of variations of the input space (e.g., different viewpoint of the same object in computer vision or paraphrasing in natural language processing tasks). Furthermore, while these workarounds may provide a temporary solution, they can be expensive, unreliable, and incompatible with the current trend of large neural models [Zhao et al. \(2023\)](#); [Chen et al. \(2022\)](#). Indeed, these large networks are typically deployed as static artifacts, whose behavior is difficult to modify during deployment without a costly full re-training [Lazaridou et al. \(2021\)](#). Thus, in all those cases, in order to adapt to changes in the environment, or to address instances of underfitting or overfitting in the original training data, it is desirable to have the ability to quickly make targeted updates to the model’s behavior after it has been deployed [De Cao et al. \(2021\)](#).

To address this need, *knowledge editing* methods have been recently proposed to efficiently change a model’s behaviors without affecting previous performance on the same task [Sinitstin et al. \(2020\)](#). These approaches take inspiration from several fields of artificial intelligence research and range from simple fine-tuning with regularization methods [Zhu et al. \(2020\)](#) to meta-learning techniques that adopt hypernetwork models to learn how to update parameters [De Cao et al. \(2021\)](#); [Mitchell et al. \(2022a\)](#). Due to its recent appearance in the literature, [Sinitstin et al. \(2020\)](#), the field still lacks accordance in the taxonomy, naming convention, datasets, and target applications. Indeed, most of the works have been motivated by large language models (LLMs), [Zhao et al. \(2023\)](#); [Brown et al. \(2020\)](#); [Soltan et al. \(2022\)](#), focusing mostly on tasks such as question answering (QA), [Levy et al. \(2017\)](#), machine translation (MT) [De Cao et al. \(2021\)](#), modifying knowledge graph embeddings [Cheng et al. \(2024\)](#), or even simpler NLP problems [Thorne et al. \(2018\)](#). However, it is also possible to find applications of knowledge editing to computer vision problems [Sinitstin et al. \(2020\)](#). Furthermore, its potential scope is poised to expand across diverse machine learning domains, encompassing areas such as medicine [Shehab et al. \(2022\)](#) robotics [Soori et al. \(2023\)](#), and precision agriculture [Sharma et al. \(2020\)](#) in the future.

1.1 Organization of the survey

The objective of this survey is to provide a comprehensive review of existing literature on knowledge editing, formalizing the task and providing a categorization of the approaches into distinct families. To our knowledge, this is the first work to undertake such an effort, and we hope that it will facilitate future research in this increasingly important area of study. Indeed, the need for more formalization and organizational structure can already be seen by a recent study [Yao et al. \(2023\)](#), which attempts to benchmark and compare some knowledge editing methodologies specifically for LLMs editing.

The rest of the survey is organized as follows. [2](#) introduces the problem of knowledge editing, using previous works to formalize it under a common definition. [3](#) explores the tasks and datasets that are most commonly considered when solving the knowledge editing problem. [4](#) provides an overview of the most relevant knowledge editing approaches in the literature, identifying four distinct families: *regularization techniques*, *meta-learning*, *direct model editing*, and *architectural strategies*. Finally, [5](#) concludes the survey by discussing some intersection with knowledge editing and other fields of research, and outlining some possible future risks and directions.

2 Overview

This section begins by presenting an introduction to the concept of *knowledge editing*, which is also referred to as *model editing* in the literature. First, we review various definitions and interpretations of knowledge editing proposed by different works. Next, we establish a common definition of knowledge editing that generalizes to all existing works in the field.

2.1 Background

The concept of knowledge editing was first introduced in [Sinitstin et al. \(2020\)](#), which formalizes it as the task of correcting a model’s mistake on a *specific sample* while preserving the model’s overall behavior, akin to continuous learning. Indeed, as specified by the authors, “The problem of efficient neural network patching differs from continual




Figure 1: The knowledge editing problem has been firstly proposed as the task of modifying a model based on a set of individual pairs of edits, in a non-sequential manner (a), [Sinitisin et al. \(2020\)](#). Successive works extended the problem to batch of edits (b), sequential individual edits (c), and sequential batch of edits (d). Evaluation metrics are similar to all cases, as described in Section 2.5.

learning, (...) [because] is not sequential in nature. However, correction (...) of mislabeled samples must not affect its behavior on other samples, which is close to overcoming [the] catastrophic forgetting task.”. Therefore, [Sinitisin et al. \(2020\)](#) define the problem as performing individual edits *reliably* and *efficiently*, not *sequentially*, and on the same task learned by the target model without affecting its overall behavior (i.e., being *local* to the edit). Concurrently, authors of [Zhu et al. \(2020\)](#) worked specifically on the task of modifying memories in Transformer models [Vaswani et al. \(2017\)](#), providing their own definition of knowledge editing. They expand the scope of the problem to a subset of knowledge, i.e., a *batch of edits*. Similarly, several other studies have also formalized the problem of model editing, following similar steps to [Zhu et al. \(2020\)](#). For instance, works by [Mitchell et al. \(2022a\)](#), [Meng et al. \(2022a\)](#), and [Mitchell et al. \(2022b\)](#) have defined the task as the process of performing individual or batch edits on a target model trained for a specific task. These studies emphasize the importance of ensuring that edits are resilient to factors of variation, i.e., that they are generalizable. While injecting individual changes is already a challenging and interesting task for the scientific community, multiple simultaneous general model edits represent a more realistic scenario that deserves further exploration.

More recently, [Hartvigsen et al. \(2022\)](#) and [Huang et al. \(2023\)](#) argued that the conventional "one-mistake-fixing scenario" does not accurately reflect the complexity of real-world knowledge editing challenges. As such, they proposed to extend the scope of knowledge editing to a sequential problem to facilitate the development of more practical editing methods. While their proposal only accounts for subsequent individual edits, considering multiple simultaneous and sequential edits represents a more general case where the number of edits varies at each step. Importantly, in the case of iterative model editing, it is desirable to respect not only the new editing task constraints but also the previous ones, which is closely related to the concept of continual learning. Nevertheless, it is crucial to highlight that while the new definition of knowledge editing acknowledges a sequential process, differing from continuous learning, its scope remains limited to the modification of knowledge of the initially learned task by the model. On the contrary, continuous learning operates without such constraints, researching for methodologies that allow the model to expand to new tasks and adapt dynamically to entirely new information.

2.2 The knowledge editing problem

To introduce the problem of knowledge editing, it is possible to leverage the terminology used to define a “well-posed learning problem” [Mitchell et al. \(2007\)](#): an algorithm is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E . Then, we can say that knowledge editing is the problem of modifying the algorithm such that, given a representative set S of instances of tasks in T , and a subset $S_e \subseteq S$ of *edit* instances, its performance on S_e improves as measured by P , while performance on all the other instances $S \setminus S_e$ remains unchanged.

More practically, we can define the problem of knowledge editing as “the task of modifying a model based on a set of individual or batch edits, S_e , pertaining to the same task known by the model, either in a sequential or non-sequential manner. The objective is to update the model’s knowledge representation without significantly altering its original behavior or performance over S and being robust to different factor of variations of the edits.”

More formally, let \mathbb{X}, \mathbb{Y} be an input and output space, respectively. Let f_0 be a function that maps an input $x \in \mathbb{X}$ to an output $y \in \mathbb{Y}$, parametrized by $\theta_0 \in \Theta$, then

$$f_0 \in \mathbb{F} := (\mathbb{X} \times \Theta)^{\mathbb{Y}} \quad (1)$$

We use the subscript zero to indicate that this is the *starting model*, i.e., the model we want to edit. We define an *edit pair* as an input-output pair $(x_e, y_e) \in \mathbb{X} \times \mathbb{Y}$, such that $f_0(x_e) \neq y_e$. Then, in its simplest form, given an individual edit example pair (x_e, y_e) , a knowledge editing (KE) methodology can be defined as follows

$$\text{KE} : \mathbb{F} \times \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{F} \quad (2)$$

i.e. a function that takes in input the starting model f_0 and the edit pair (x_e, y_e) to produce an *edited model* f_e . If the edited model is such that $f_e(x_e) = y_e$, then we say that the edit was successful. A KE approach can realize the transformation from f_0 to f_e in different ways, and we identify four families of possible realizations: regularization, meta-learning, direct model editing and architectural strategies. We deep dive into more details into each of these families in 4.

2.3 Taxonomy of edits

Often, it is interesting to edit a model applying multiple edits at once, a sequence of edits, or sequences of batches of edits.

The definition provided in section 2.2 has been given for the simplest case, that we call a *single non-successive* edit: we only want to change the model for one edit pair. Conversely, for multiple successive edits we can formally define a list of edit requests as:

$$\mathcal{E} = \{(x_e, y_e)^{(i)} \text{ s.t. } \forall i, j \ x_e^i = x_e^j \Rightarrow y_e^i = y_e^j\} \quad (3)$$

where the logical constraint ensures that there are no conflicting requests, as suggested by Meng et al. (2022a). Individual edits can also be grouped together and form N batches of successive edits each with $B_e^{(i)}$ edit pairs, such as

$$\mathcal{B}_e^{(i)} = \{(x_e, y_e)^0, \dots, (x_e, y_e)^B\}^{(i)} \text{ s.t. } \mathcal{E} = \bigcup_{i=1}^N \mathcal{B}_e^{(i)} \quad (4)$$

The difference between successive individual or batch of edits is that some KE methodologies can ingest an entire $\mathcal{B}_e^{(i)}$ and produce the resulting f_e implementing all the given edits, while other methodologies can only consider one individual sample in sequence at a time. In both cases (a sequence of individual edits is trivially a sequence of single-item batch edits), successive edits assume to work with a starting model f_{t-1} and apply the t -th change as

$$f_t = \text{KE}(f_{t-1}, \mathcal{B}_e^{(i)}) \quad (5)$$

proceeding iteratively, using f_t as a starting model for the next edit. Figure 1 summarizes the four types of edits. Finally, as for individual edits, after every change, f_e should not only satisfy $f_e(x_e) = y_e$, but a series of other properties as discussed in the next section.

2.4 Editing properties

Based on the specific task learned by function f , various properties can be specifically designed. However, at a fundamental level, following Sinitin et al. (2020); Huang et al. (2023), knowledge editing should aim at satisfying four properties, that we define below.

Property 1 - Reliability: Given an edit pair (x_e, y_e) , the edited model f_e should output the desired edit:

$$f_e(x_e) = y_e \quad (6)$$

Property 2 - Generality: The edited model f_e should be able to generalize to *similar* examples to the edit pair. This can be formalized by defining an *equivalence neighborhood* $N(x_e) \subset \mathbb{X}$ and requiring that the edited model f_e satisfies:

$$f_e(\tilde{x}_e) = y_e \quad \forall x_e \in N(x_e) \quad (7)$$

Property 3 - Locality: The edited model f_e should not alter the output of examples that are not similar to the edit pair. This can be formalized by defining a *locality set*

$$L(x_e) = \{(x_{loc}, y_{loc}) \in \mathbb{X} \times \mathbb{Y} \mid \text{s.t. } x_{loc} \notin N(x_e) \wedge f_0(x_{loc}) = y_{loc}\} \quad (8)$$

and require that the edited model f_e satisfies:

$$f_e(x_{loc}) = y_{loc} \quad \forall (x_{loc}, y_{loc}) \in L(x_e) \quad (9)$$

Property 4 - Efficiency: The proposed knowledge editing methodology should aim to achieve efficient editing, characterized by low compute and space complexity.

It is worth noting that some works in the literature have defined these same properties, or a subset of them, with different names and notations. Specifically, reliability may also be referred to as *efficacy*, locality as *specificity*, and generality as *paraphrase*, Meng et al. (2022b); De Cao et al. (2021). Moreover, as previously mentioned, depending on the specific field of application, other properties and related evaluation metrics can arise, such as *fluency* and *consistency* Meng et al. (2022b) when editing Language Models. Finally, the last property, namely efficiency, tends to be disregarded in academic literature. However, it is one of the main reasons KE is appealing over a simpler re-training of the base model. Furthermore, it plays a pivotal role in performing comparative analyses with baseline models and showcasing the ability of a particular methodology to modify a neural network with a substantial number of parameters. We encourage future works to consider (at least) all four properties when measuring the effectiveness of a proposed KE methodology.

2.5 Evaluation metrics

Given a knowledge editing algorithm KE, a list of edit requests \mathcal{E} , a starting model f_0 and a test set \mathcal{D}_{test} , it is possible to measure the extent to which the properties above hold. We consider working with N batches of successive edits, each comprised of B individual edits, without loss of generality (as mentioned above, if $N = 1$, we have successive individual edits). Moreover, if the KE methodology is unable to process all B edits concurrently, it can apply each edit individually in a non-sequential manner.

Following Huang et al. (2023), we represent with I the indicator function, and define the following metrics.

Success Rate (SR): It is used to evaluate reliability, and it is simply the proportion of edits for which the methodology succeeds in changing the knowledge of a starting model f_t .

$$\begin{aligned} \text{SR} &= \frac{1}{N} \frac{1}{B} \sum_{n=1}^N \sum_{b=1}^B I(f_{n,B}(x_{e;n,b}) = y_{e;n,b}) \\ \text{s.t. } &\bigcup_{n=1}^N \bigcup_{b=1}^B (x_e, y_e)_{n,b} = \mathcal{E} \end{aligned} \quad (10)$$

In the case of non-sequential individual edits, $f_{n,B} = f_{n,b}$. Moreover, Eq. 10 provides an overall value after N successive edits, but it can be of interest to measure SR every n edits, tracking changes over the sequence.

Generalization Rate (GR): It is used to evaluate generality, testing the post-edit model f_e , on the equivalence neighborhood set $N(x_{e;n,b}, y_{e;n,b})$, where $(x_{e;n,b}, y_{e;n,b})$ is the n -th batch, b -th edit pair. GR can be written as,

$$\begin{aligned} \text{GR} &= \frac{1}{N} \frac{1}{B} \frac{1}{N_b} \sum_{n=1}^N \sum_{b=1}^B \sum_{i=1}^{N_b} I(f_{n,B}(\tilde{x}_{e;n,b,i}) = \tilde{y}_{e;n,b,i}) \\ \text{s.t. } &\forall n, b \bigcup_{i=1}^{N_b} (\tilde{x}_e, \tilde{y}_e)_{n,b,i} \subseteq N(x_e, y_e)_{n,b} \end{aligned} \quad (11)$$

where N_b is the number of equivalent samples of the b -th edit pair. Following Mitchell et al. (2022a), we can also define **Edit Success (ES)** to summarize both SR and GR. It can be computed as the average accuracy of the edited model f_e on the edit input(s), as well as inputs drawn from the equivalence neighborhood(s), that is,

$$\begin{aligned} \text{ES} &= \frac{1}{N} \frac{1}{B} \sum_{n=1}^N \sum_{b=1}^B \left(I(f_{n,B}(x_{e;n,b}) = y_{e;n,b}) \right. \\ &\quad \left. + \sum_{i=1}^{N_b} \frac{1}{N_b} I(f_{n,B}(x_{e;n,b,i}) = y_{e;n,b,i}) \right) \end{aligned} \quad (12)$$

where the same conditions for SR and GR hold and have been omitted for brevity.

Drawdown (DD): It is used to evaluate locality, and it is defined as the performance degradation of the edited model over \mathcal{D}_{test} . It is computed using the final edited model, f_N , that in case of successive edits is the result of N steps.

$$DD = 1 - \frac{\sum_{(x,y) \in \mathcal{D}_{test}} I(f_N(x) = y)}{\sum_{(x,y) \in \mathcal{D}_{test}} I(f_0(x) = y)} \quad (13)$$

Finally, as suggested by [Huang et al. \(2023\)](#) in the case of multiple successive edits, it is also important to evaluate SR and GR using the final model f_N , in order to assess how past edits are retained. Therefore, it is possible to define three additional metrics, **Success Retain Rate (SRR)**, **Generalization Retain Rate (GRR)**, and **Edit Success Retain (ESR)**, simply using in Eq. 10 and 11, f_N instead of $f_{n,B}$.

3 Tasks and Datasets

The formalization of the knowledge editing problem provided in the previous section is general, and many applications of knowledge editing to different tasks encompassing various fields can be formulated within that framework. The brief, but rich history of the field has so far seen applications mainly to two broad fields: Computer Vision and Natural Language Processing. Indeed, [Sinitin et al. \(2020\)](#) provides experimental results on image classification and machine translation, and almost all the works that come after (and even before [Garnelo et al. \(2018\)](#); [Kirkpatrick et al. \(2017\)](#)) demonstrate the effectiveness of their proposed approaches in one or more applications in these two fields.

Even though the knowledge editing framework can be defined independently of the target domain and task, each specific application has its unique challenges and intricacies, which we explore in this section. 3.1 covers the most common tasks in the Computer Vision domain, as well as the datasets on which the tasks are usually addressed, and how the knowledge editing problem can be instantiated in this context. Section 3.2 provides a similar overview for applications to the Natural Language Processing domain. Finally, 3.3 describes tasks and datasets that do not strictly fit in any of the two domains above.

3.1 Computer Vision

Computer Vision is a broad field with a long history, which generally attempts to extract meaningful representations from visual media to derive an understanding of the represented scenes [Szeliski \(2022\)](#). Over the last years, deep learning methods have been shown to outperform previous state-of-the-art techniques in several applications in the field [Voulodimos et al. \(2018\)](#), and while more “traditional” techniques are still relevant for some applications, neural networks-based approaches have become the de facto standard for many others [O’Mahony et al. \(2020\)](#). Due to the importance and breadth of the field, and the relevance of neural networks therein, knowledge editing literature has found fertile grounds in Computer Vision, and has so far gravitated towards two primary tasks: Image Classification and Image Completion. A number of datasets are customarily used to test approaches to solve these tasks. They vary in terms of number of examples, classes and channels, as well as resolution of the representative images; Table 1 provides an overview of the most commonly used ones.

Image Classification The task of image classification is straightforward, we wish to label a complete image (or predetermined portion) with its most likely semantic category, e.g., horse, cat, or car [Szeliski \(2022\)](#). In this context, an example is an image and its semantic label. The image, of predefined dimension (or resolution), is encoded as a 3D tensor $x^{(i)} \in \mathbb{R}^{W \times H \times C}$, where W is the width of the image, H the height of the image, and C the number of channels, depending on whether the image is grayscale (1) or RGB (3) or something else (e.g., RGBD [Firman \(2016\)](#) with an additional depth channel). The editing task is then often formulated by artificially corrupting a subset of either the images or the labels. The latter is the more prevalent, and usually involves training models on a subset of a dataset and subsequently introducing random label shuffling within a withheld set to create an edit set $(x_e^{(i)}, y_e^{(i)})_{i=1}^N$, where originally $y^{(i)} \neq y_e^{(i)}$. Other works, such as [Sotoudeh and Thakur \(2021\)](#), introduce e.g. motion blur or fog [Mu and Gilmer \(2019\)](#) to corrupt a subset of the images, creating an edit set where originally $x^{(i)} \neq x_e^{(i)}$.

Several datasets support knowledge editing experiments for this task, and a distinction is often made among “toy” and “large-scale” datasets. Well-known datasets like MNIST [LeCun et al. \(2010\)](#) and CIFAR-10 [Krizhevsky \(2009\)](#), which are widely used in the literature, are frequently employed for experimentation and belong to the former category. For more challenging and realistic scenarios, researchers turn to larger scale datasets, of which the most popular is surely the extensive ImageNet Database [Deng et al. \(2009a\)](#), which now encompasses over 10 million labeled images, spanning more than 20,000 object categories. Specifically, studies such as [Sinitin et al. \(2020\)](#) and [Lee et al. \(2019\)](#) explore

| Dataset | Tasks | # Examples | # Classes | # Channels | Resolution |
|-----------|----------------------------|------------|-----------|------------|------------|
| MNIST | Classification, Inpainting | 70k | 10 | 1 | 28x28 |
| CIFAR-10 | Classification | 60k | 10 | 3 | 32x32 |
| CIFAR-100 | Classification | 600k | 100 | 3 | 32x32 |
| ImageNet | Classification | 1.2 M | 1000+ | 3 | 224x224 + |
| NAE | Classification | 7.5k | 200+ | 3 | 224x224 + |
| CelebA | Inpainting | 200k | - | 3 | 178x218 |

Table 1: Most important datasets used in Computer Vision for Knowledge editing. MNIST, CIFAR-10, CIFAR-100 are generally regarded as “toy” datasets while ImageNet, NAE, CelebA as more challenging testbeds.

datasets derived from the ImageNet Large Scale Visual Recognition Challenges (ILSVRC) [Russakovsky et al. \(2015\)](#). To further accentuate the complexity, [Sinitin et al. \(2020\)](#) introduces a highly challenging configuration, leveraging the Natural Adversarial Examples (NAE) dataset [Hendrycks et al. \(2021\)](#), consisting of 7500 natural images notorious for their arduous classification nature, where pre-trained models exhibit correct prediction rates of less than 1% for NAEs.

Image Inpainting Image inpainting, also known as image completion, is the task of reconstructing missing regions in an image [Szeliski \(2022\)](#). The problem is formalized as a regression over functions mapping pixel coordinates within $[0, 1]^2$ to pixel values in $[0, 1]$ (grayscale) or $[0, 1]^3$ (RGB). This task has so far received less attention from the knowledge editing community [Garnelo et al. \(2018\)](#), leveraging datasets that encompass both the MNIST dataset, once again serving as a rudimentary example, and the CelebFaces Attributes Dataset (CelebA) [Liu et al. \(2015\)](#) for more challenging scenarios. The CelebA dataset presents a more demanding scenario, offering over 200,000 celebrity images, each accompanied by 40 attribute annotations, making it a challenging and comprehensive dataset for exploration.

3.2 Natural Language Processing

Natural Language Processing (NLP) is also a broad field, concerned with giving computers the ability to process and understand human language [Eisenstein \(2019\)](#). Like in computer vision, in recent years researchers and practitioners in the field have leveraged the power of neural networks with many outstanding results [Otter et al. \(2020\)](#); [Soltan et al. \(2022\)](#); [FitzGerald et al. \(2022\)](#). With the recent paradigm shift from supervised learning to pre-training followed by fine-tuning [Wang et al. \(2022\)](#), and the trend towards larger and larger models [Zhao et al. \(2023\)](#), the ability to perform a cheap edit of a model instead of an expensive fine-tuning has motivated an intense interest from the knowledge editing community. Undoubtedly, within the NLP realm, the most widely targeted tasks for knowledge editing are Fact-Checking when dealing with classification, and (closed book) Question Answering for language generation. Some recent works also explore open Text Generation, editing of factual relations with had hoc datasets and also Document Classification. Table 2 provides an overview of the datasets commonly used for those tasks.

Fact-checking Fact-checking is the task of assessing whether claims made in written or spoken language are true, often addressed as a binary classification task [Guo et al. \(2022\)](#). In this setting, examples are natural language claims coupled with binary labels, even though occasionally a third neutral option is available. The claim $x^{(i)}$ is encoded with some tokenization scheme (or pre-trained word embedding) as a sequence of integers (or semantic vectors), while the label $y^{(i)}$ can take one of two values, positive or negative (optionally a third, aforementioned neutral value). One can then have a neural network predict this value with an explicit classification layer, or alternatively a language model producing special tokens (e.g., *True/False* or *Supports/Refutes*) for the considered truth values, when prompted with the claim under consideration. The Fact Checking task has been considered particularly appealing for the knowledge editing community [De Cao et al. \(2021\)](#); [Mitchell et al. \(2022a,b\)](#); [Huang et al. \(2023\)](#) for at least a couple of reasons. First, the recent phenomenal success of language models also highlighted their proneness to generate reasonable but factually wrong natural language text [Ortega et al. \(2021\)](#); [Ji et al. \(2023\)](#). This degrades the system performance and fails to meet user expectations in many real-world scenarios, leading to a great interest in the ability to mitigate these hallucinations. Furthermore, reasonable edit sets are fairly easy to obtain, e.g. randomly flipping the labels of claims from pre-existing datasets from Support to Refutes and vice versa. The most widely used datasets for this task are FEVER [Thorne et al. \(2018\)](#) and VitaminC [Schuster et al. \(2021\)](#). Both of them are extracted from Wikipedia and annotated by human experts, to arrive at (*evidence, wikipage, claim, label*) tuples. In both cases, the label can be *Supports, Refutes* or *Not Enough Info* depending on whether the evidence supports or not the claim. To construct proper editing datasets $\{(x_e^{(i)}, y_e^{(i)})_{i=1}^N\}$ out of them, [De Cao et al. \(2021\)](#) (for FEVER) and [Mitchell et al. \(2022b\)](#)

(for VitaminC) grouped facts based on the same pages, augmented each fact $x_e^{(i)}$ with some rephrases $\tilde{x}_e^{(i)}$, to assess generality, and randomly flipped the label $y_e^{(i)} \neq y^{(i)}$ in each group.

Question Answering The task of training a model to provide a correct natural language answer to a given question is referred to as question answering; more specifically, *closed-book* question answering is restricted to the case when the model is only fed with the question, and not a selection of possible answers, or supporting corpus from which an answer can be extracted Roberts et al. (2020). In this case, both the input x and the target label y are natural language sequences, and we test the extent to which the model parameters implicitly store the required knowledge to provide correct answers. The most widely adopted dataset is surely the zero-shot Relation Extraction dataset (zsRE) Levy et al. (2017), used in particular by Zhu et al. (2020); De Cao et al. (2021); Mitchell et al. (2022a,b); Meng et al. (2022a,b); Huang et al. (2023); Hartvigsen et al. (2022). This task and datasets are particularly appealing for knowledge editing. Indeed, *closed-book* question answering can benefit greatly from knowledge editing, as pipelines that solve the task usually leverage factual knowledge to answer questions; in the case of neural networks, this knowledge is acquired during training and implicitly stored in the networks’ parameters, and it is unclear how to tweak these parameters to correct wrong or undesired answers, especially as the networks grow bigger in size. Meng et al. (2022b) hypothesize that this factual knowledge takes the form of (*relation, subject, object*) triples, with intermediate layers acting as key-value storage units. This formalization lends itself nicely to the definition of an editing objective, rather than directly the open-ended natural language generation task. Furthermore, Levy et al. (2017) demonstrates that it is possible to reduce relation extraction to the problem of answering simple reading comprehension questions and provided in their dataset multiple templates for each relation. For example, the triple (*occupation, s, o*) can be naturally extracted by answering one of the following questions: *What did s do for a living?*, *What is s’s job?* *What is the profession of s?*. The subject s can then be modified to create editing examples. Section 4.3 further discusses factual knowledge and how different works have modeled it for improving knowledge editing. Besides zsRE, knowledge editing of models solving Question Answering has been studied leveraging also on additional datasets such as t-REX Elshahar et al. (2018) and Natural Questions (NQ) Kwiatkowski et al. (2019). Finally, as a more challenging flavor of the same task with added counterfactual information Meng et al. (2022b) introduced a new dataset called Counterfact.

Further NLP tasks Beside the two popular tasks outlined above, Mitchell et al. (2022a) tested editing text generation from autoregressive GTP-like models on a special version of WikiText-103 Merity et al. (2016), where they are considering as prompts (x_e) passages sampled from WikiText itself and as edit targets (y_e) 10-token samples from a pre-trained distilGPT-2 model. This was for them a valid challenging editing setup since for their target model for editing a greedy 10-token prediction agrees with these edit targets for < 1% of examples they extracted. Finally, more recently, Hartvigsen et al. (2022) tested their methodology on a novel task for knowledge editing using the SCOTUS dataset from Chalkidis et al. (2022). The classification task is to categorize U.S. Supreme Court documents over multiple decades into 11 topics. What makes this task interesting is that, over time, categorization rules change, so that label distributions shift. We note how this is off the shelf particularly realistic for knowledge editing as much of the world knowledge memorized by networks evolves over time just like those labels shifts in the dataset and the target of knowledge editing can be seen as keeping updated such world knowledge.

| Dataset | Tasks | Format | # Examples | # Classes |
|-------------|-------------------------|--|------------|-----------|
| FEVER | Fact Checking | (evidence, wikipage, claim, label) | 420k | 3 |
| VitaminC | Fact Checking | (evidence, wikipage, claim, label) | 450k | 3 |
| zsRE | Question Answering | (subject, relation, object) | 120M | - |
| T-REx | Question Answering | (subject, relation, object) | 11M | - |
| NQ | Question Answering | (question, answer) | 320k | - |
| CounterFact | Question Answering | (subject, relation, true object, false object) | 22k | - |
| Wikitext | Text Generation | tokens | 100M | - |
| SCOTUS | Document Classification | (date, text, label) | 9.2k | 11 |

Table 2: Most important datasets for knowledge editing used in NLP. We report characteristics of the original datasets even though often for knowledge editing ad hoc version, preprocessed to make editing meaningful are used.

3.3 Other Applications

Even though the majority of works in the knowledge editing literature has focused on the Computer Vision and Natural Language Processing fields, as described above, the general nature of the editing problem yielded interesting results also in other fields, and will likely yield more in more diverse fields and applications in the years to come. Among

these, to the best of our knowledge, the most notable examples are applications in safety-critical scenarios and to graph neural networks; in the following we briefly review works from both.

Safety-critical Systems Safety-critical systems are those systems whose failure may lead to consequences that are determined to be unacceptable, such as significant damage to properties, the environment, or people Knight (2002). Deep neural networks have grown in popularity over the past decade and are now being used in safety-critical domains such as self-driving cars Gupta et al. (2021), healthcare Tekkeşin et al. (2019) and aviation Sridhar (2020). Clearly, in such critical scenarios, being able to find and correct unsafe neural network behavior becomes a crucial objective. This has motivated a line of research within the knowledge editing community, that so far has only touched the aviation domain, specifically the aircraft collision avoidance problem. The systems currently employed provide safety guarantees at the cost of being poorly data efficient, and efforts have been made to integrate neural networks into the pipeline Julian et al. (2016); Julian and Kochenderfer (2019). As a consequence, several subsequent works Sotoudeh and Thakur (2021); Fu and Li (2021); Liang et al. (2023) from the knowledge editing community have proposed approaches for fixing unsafe behavior of neural networks integrated in safety-critical pipelines.

Developing a robust collision avoidance algorithm that reliably prevents collision without alerting excessively is challenging due to sensor error and uncertainty in the future behavior of the aircraft. The Airborne Collision Avoidance System X (ACAS X) Kochenderfer and Chryssanthacopoulos (2011) family of collision avoidance systems formulates the problem of collision avoidance as a partially observable Markov decision process. The variant for unmanned aircraft, ACAS Xu, uses dynamic programming (DP) to then find a solution in terms of resolution advisories that avoid collisions while minimizing disruptive alerts. The DP process makes use of a massive lookup table, that makes storage costly and certification time-consuming for certified avionics systems. Therefore, Julian et al. (2016) propose using a deep neural network for compressing the table without loss of performance, as measured by a set of safety and operational metrics. There are seven real-valued state variables that define an aircraft encounter, describing its geometry in terms of the two aircraft involved: (1) distance from ownship to intruder (2) angle to intruder relative to ownship heading direction (3) heading angle of intruder relative to ownship heading direction (4) speed of ownship (5) speed of intruder (6) time until loss of vertical separation (7) previous advisory action. There are then five possible horizontal maneuver *advisories* that the system can produce: clear-of-conflict, or adjusting course by turning left or right at two fixed angles (hence 4 more possibilities). The state variables are usually discretized arriving at ≈ 120 million points, and the aforementioned lookup table associates scores to all pairs of 120 million states and five actions. This table is what makes up the ACAS Xu dataset: $\{(x^{(i)}, y^{(i)})_{i=1}^N\}$ with $N = 5 \times 120$ million, where $x^{(i)}$ represents a discretized seven-dimensional state, and $y^{(i)} \in \mathbb{R}^5$ is the vector of scores associated to each of the five possible actions in that state. With 600 million floating-point numbers, the table requires over 2 GB of storage. The task for the neural network is to regress this table, minimizing parametric knowledge required (i.e., number of parameters) and error with respect to the table. It is interesting to note that this is an atypical regression problem, since we aim for the guarantee that the optimal advisory remains the same. When the difference between the scores of the first and second-best advisories is relatively small, simple regression techniques (e.g., minimizing the Mean Squared Error) can lead to the network realizing a different strategy from that of the original table. This is reflected in the design of the loss function. The network or its further refinements Julian and Kochenderfer (2019) is then verified via tools such as Wang et al. (2018); Katz et al. (2017), that are able to prove several input-output-based security properties, e.g., that a clear-of-conflict advisory will always be issued if the intruder is sufficiently far away, thus providing formal guarantees about DNN behavior. These properties are formalized as implications of the form $\forall x, x \in B \implies f_\theta(x) \in C$, where f_θ is the function approximator realized by the DNN with parameters θ , B is a bounded region of the input space and C a bounded region of the output space. The ACAS Xu case study has so far been of great interest for the knowledge editing community since one such security property has been found to not be satisfied by the original network, exposing an input on which the network was inconsistent with the lookup table (ϕ_8 in Katz et al. (2017)). This discrepancy would then be addressed by retraining the DNN, thus leading to the central question of knowledge editing: how to fix the network behavior on a limited set of points without affecting its behavior on unrelated points. Each work mentioned at the beginning of the paragraph has addressed this problem differently, but sharing the same setup: once a failing security property for a network is identified, and one is able to generate counter-examples, i.e., pairs $(x^{(i)}, y^{(i)})$ such that $x^{(i)} \in B$ but $y^{(i)} \notin C$, a certain strategy for defining candidate edit pairs can be formalized, i.e., how to assign a particular $\bar{y}^{(i)}$ to $y^{(i)}$. Then, usually a subset of this becomes the edit set, while the remaining portion is chosen as generality set; finally, a locality set is defined by points correctly classified by the network, i.e., input-output pairs for which the properties under scrutiny hold true.

Graph Neural Networks Deep learning models have been particularly successful when dealing with signals such as speech, images, or video, in which there is an underlying Euclidean structure; however, recently, there has been a growing interest in trying to apply learning on non-Euclidean geometric data, for instance represented in the form of graphs Bronstein et al. (2017). Graph Neural Networks (GNNs) learn node representations by applying shared

permutation invariant functions over local neighborhoods of nodes in the graph Bronstein et al. (2021). These representations can then be used for tasks like node classification. For instance, assigning a category to each paper in a citation graph Wu et al. (2020). GNNs have achieved prominent results in learning features and topology of graph data, however, knowledge editing for GNNs is rarely explored, despite their widespread applicability; therefore, Liu et al. (2023) propose a method to edit these models, restricting to the aforementioned node classification task.

The task can be formalized as follows: let $G = (V, E)$ be an undirected graph with $V = (v_1, \dots, v_{|V|})$ and $E = (e_1, \dots, e_{|E|})$ being the set of nodes and edges, respectively. Given a feature space \mathcal{X} (e.g., the space of real-valued d -dimensional feature vectors \mathbb{R}^d), a node feature tensor $X \in \mathcal{X}^{|V|}$ and a label space \mathcal{Y} , the goal is to learn a representation h_v from which a label $y_v \in \mathcal{Y}$ for each node $v \in V$ can be easily predicted. Many datasets for node classification exist in the literature, comprised of data from various domains like citation networks and social networks, and we find again the distinction between small-scale and large-scale datasets Wu et al. (2020). Among the former we find datasets like Cora, which contains a selection of Machine Learning papers collected from the web, and the references automatically parsed from their bibliographies McCallum et al. (2000). In particular, the network contains $|V| = 2708$ nodes (articles) and $|E| = 5429$ links (citations); the feature space is $\mathcal{X} = \{0, 1\}^{1433}$, i.e., each node is described in terms of the presence (or absence) of certain keywords, taken from a dictionary of 1433 unique words (bag-of-words content representation); the label space is $\mathcal{Y} = \{1, \dots, 7\}$, i.e. the task is to predict to which of seven classes (e.g., *Theory* or *Reinforcement_Learning*) each publication belongs. The Reddit dataset is instead a popular representative of large-scale node classification datasets. Hamilton et al. (2017). constructed a graph dataset from Reddit posts made in the month of September 2014. The node label $y_v \in \mathcal{Y}$ in this case is the community, or “subreddit”, that a post belongs to, considering $|\mathcal{Y}| = 50$ large communities. A post-to-post graph is constructed connecting posts if the same user comments on both. In total, this dataset contains $|V| = 232,965$ posts with an average degree of 492 ($|E|$ is in the order of 100 million edges). In both of these cases, and the many other instances of node classification datasets, constructing an edit dataset is fairly straightforward, and done in the same manner as for the image classification task in Computer Vision. After training the model $f_0(\cdot)$ under consideration on a *subgraph*, one evaluates it on the whole graph: each pair $(x^{(i)}, y^{(i)})$ such that $\hat{y}^{(i)} = \arg \max f(x^{(i)})$ is incorrect, becomes an edit pair $(x_e^{(i)}, y_e^{(i)})$. The geometry of graphs lends itself nicely to defining also generality and locality sets: indeed, since the task under consideration is node classification, as we have seen a single example $x^{(i)}$ describes a node v , then one can define its neighborhood $N(x^{(i)})$ to be its actual neighborhood in the graph $N_G(v) = \{w \in V \mid \exists (v, w) \in E\}$; from this definition, generality, and locality sets follow consequently, as seen in earlier sections.

4 Knowledge Editing Methodologies

In recent times, several “knowledge editing” methods have been introduced to effectively modify the behaviors of models while maintaining their previous performance on the same task Sinitsin et al. (2020). These approaches draw inspiration from various fields of artificial intelligence research and can be broadly categorized into four distinct families: *regularization techniques*, *meta-learning*, *direct model editing*, and *architectural strategies*.

Regularization techniques utilize various forms of regularization to guide the model’s learning process during fine-tuning, encouraging it to incorporate the desired edits while retaining its original capabilities Zhu et al. (2020). Meta-learning approaches, on the other hand, employ hypernetwork models to learn parameter updates, enabling efficient adaptation to new tasks or knowledge De Cao et al. (2021); Mitchell et al. (2022a). Direct model editing methods involve directly modifying the model’s parameters or representations to incorporate the desired changes. These techniques can range from simple parameter updates to more complex approaches that leverage the model’s internal representations. Finally, architectural strategies explore modifying the model’s architecture itself, either by introducing new components or restructuring existing ones, to facilitate the integration of new knowledge or behaviors. In the upcoming sections, we will provide detailed discussions on each of these families of knowledge editing methodologies, highlighting their respective areas of application, advantages, limitations, and notable techniques within each category.

The objective of this section is to categorize various knowledge editing techniques discussed in the literature into the four distinct families mentioned above. All presented works have different characteristics, target different areas of application, types of edits, and adopt diverse experimentation strategies. Nevertheless, the objective of all the works reported can be formulated within the formal framework given in 2.2. A comparison between the most notable KE methodologies at the time of writing can be found in Table 3. On the other hand, Table 4 presents a comparison with non-sequential single batch edits of factual knowledge.

| KE Methodology | KE Category | Training Required | Preserves Architecture | Only (x_e, y_e) | SNS | BNS | SSE | BSE | Scales to LM |
|-----------------|----------------|-------------------|------------------------|-------------------|-----|-----|-----|-----|--------------|
| FT + L_2 | Regularization | False | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✗ |
| FT + KL | Regularization | False | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✗ |
| EWC | Regularization | False | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✗ |
| CNP | Architectural | True | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| ENN | Meta-Learning | True | ✗ | ✓ | ✓ | ✗ | ✓ | ✗ | ✗ |
| KnowledgeEditor | Meta-Learning | True | ✓ | ✓ | ✓ | ✗ | ✗ | ✗ | ✓ |
| MEND | Meta-Learning | True | ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | ✓ |
| MALMEN | Meta-Learning | True | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ |
| ROME | Direct Editing | False | ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | ✓ |
| MEMIT | Direct Editing | False | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ |
| PMET | Direct Editing | False | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ |
| SERAC | Architectural | True | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| CaliNet | Architectural | True | ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | ✓ |
| T-Patcher | Architectural | False | ✗ | ✗ | ✓ | ✓ | ✓ | ✗ | ✓ |
| GRACE | Architectural | True | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 3: Comparison of the most notable KE methodologies in the literature. Different characteristics are reported for each approach, highlighting main advantages and disadvantages. For all approaches, we report: the category and whether it requires *training* of an auxiliary model; if it *preserves the architecture* of the edited model or requires the introduction of new components; whether it needs only the edit pair (x_e, y_e) , or requires additional input to perform the edit; if it is able to handle *single non-successive* edits (SNS), *batched non-successive* edits (BNS), *single successive* edits (SSE), and *batched successive* edits. Finally, we report if it can scale to Large Models (LM), that, following the definition in [Zhao et al. \(2023\)](#), are models with more than 10B parameters.

4.1 Regularization Techniques

Catastrophic forgetting, [Kemker et al. \(2018\)](#), is a well-known phenomenon in literature, fundamentally limiting the flexibility of editing networks once trained or fine-tuned, [Lee et al. \(2019\)](#). Indeed, in the absence of any regularization process, the regular fine-tuning signal has the ability to easily execute a specific edit, albeit with a tendency to over-fit to the provided edit examples. However, this approach lacks in providing generality to the edit and has a negative impact on locality, owing to the absence of a plasticity mechanism, [Simitsin et al. \(2020\)](#). Similar to continual learning, regularization techniques for knowledge editing aim to modify the standard fine-tuning signal of the target edit to ensure reliability and locality. Therefore, for regularization techniques, KE is not parametrized and does not require any pre-training, but it is nothing more than gradient descent computed with given edits and some specific regularization terms. While not all of these techniques were originally developed for the specific task of knowledge editing, they have proven to be in some way effective and are commonly used as useful baselines for comparison. Moreover, due to their simplicity, they can easily adapt to work with different types of edits: from single non-successive edits to batches of successive edits. However, depending on the number of layers fine-tuned, they rarely scale to models with large number of parameters such as Large Language Models (LLMs), that, according to [Zhao et al. \(2023\)](#) are models with more than 10B parameters. In these cases, methodologies discussed in the following sections, may be better suited for efficiently editing large-scale with constrained resources.

For instance, authors in [Sotoudeh and Thakur \(2021\)](#) introduce a technique called “Neural Network Patching” that allows for the correction of deep neural networks without retraining from scratch. They present the concept of neural network patching, which involves identifying and correcting the faulty or damaged components of a network. The proposed method involves using a set of repair operators to identify and replace damaged components of the network, identifying the minimum L_2 norm parameter update that reliably edits the model output, while minimizing the deviation from the original target model. Conversely, [Zhu et al. \(2020\)](#) utilize a constrained optimization process instead of penalizing an updated model for deviating from the original one. Their method employs either an L_2 or L_∞ constraint between the original model parameters and the edited ones, highlighting the importance of selecting a subset of parameters to be updated. Additionally, they demonstrate how their method is similar to Elastic Weight Consolidation (EWC), [Lee et al. \(2019\)](#), which involves computing a penalty term for each weight in the neural network based on how much it contributes to the performance of the original task. Similarly to [Zhu et al. \(2020\)](#), loss enriched with the Kullback–Leibler divergence (KL) have been proposed in order to regularize the network’s weights based on the KL divergence between the network’s output on the old task and the network’s output on the new task. This encourages the

| Dataset | Model | Evaluation Metrics | FT | KE | MEND | ROME | MEMIT | SERAC | T-Patcher |
|-------------|------------|--------------------|-------|-------|-------|-------|-------|-------|-----------|
| zsRE | GPT-J [6B] | SR | 54.70 | 6.60 | 45.60 | 99.18 | 99.23 | 90.16 | 97.12 |
| | | GR | 49.20 | 7.80 | 48.00 | 94.90 | 87.16 | 89.96 | 94.95 |
| | | DD | 62.76 | 5.82 | 11.79 | 0 | 0 | 0.1 | 3.76 |
| CounterFact | GPT-J [6B] | SR | 99.90 | 13.40 | 73.80 | 99.80 | 99.90 | 99.78 | 100.00 |
| | | GR | 97.53 | 11.00 | 74.20 | 86.63 | 73.13 | 99.41 | 83.98 |
| | | DD | 98.98 | 5.62 | 96.25 | 0 | 0 | 1.11 | 91.63 |

Table 4: Comparison of some of the most notable KE methodologies on zsRE [Levy et al. \(2017\)](#) and CounterFact [Meng et al. \(2022b\)](#) with non-sequential single batch edits of factual knowledge. For each methodology, we report success rate (SR), generalization rate (GR), and Drawdown (DD) metrics. Results adapted from [Yao et al. \(2023\)](#).

network to maintain similar weights for the parts of the network that are relevant to both tasks [Yoon et al. \(2017\)](#); [Serra et al. \(2018\)](#); [Mitchell et al. \(2022a\)](#); [Huang et al. \(2023\)](#).

As previously mentioned, the techniques outlined in this section have been primarily utilized as baseline approaches in knowledge editing and heavily intersect with continual learning works [Mundt et al. \(2023\)](#). However, in various experiments cited in [De Cao et al. \(2021\)](#); [Mitchell et al. \(2022a\)](#); [Huang et al. \(2023\)](#), these techniques have demonstrated limited efficacy in maintaining model accuracy with regard to previous knowledge. In fact, as noted by [De Cao et al. \(2021\)](#), regularization techniques overlook the highly nonlinear nature of large models and the significant impact that even minor changes in parameter space can have on the output of numerous data points. For those reasons, more elaborated works have been proposed in literature as the ones in the following sections.

4.2 Meta-Learning and HyperNetworks

"Meta-learning techniques refer to a set of algorithms and approaches that enable machines to learn how to learn [Finn et al. \(2017\)](#). These techniques have proven to be particularly useful in scenarios where a model needs to adapt quickly to new tasks or environments with limited data. For instance, in computer vision [Ren et al. \(2018\)](#), meta-learning can be employed to rapidly fine-tune a pre-trained model to recognize new object categories or adapt to different imaging conditions. Similarly, in natural language processing [Gu et al. \(2020\)](#), meta-learning can facilitate the adaptation of language models to different domains or styles of text with minimal additional training data. Robotics [Finn et al. \(2017\)](#) is another domain where meta-learning has shown promise, enabling robots to rapidly acquire new skills or adapt to changes in their environment. In reinforcement learning [Houthoofd et al. \(2018\)](#), meta-learning can help agents generalize their learned policies to new environments or task variations, improving their sample efficiency and adaptability.

The key advantage of meta-learning techniques lies in their ability to leverage prior knowledge and experience to learn new tasks or adapt to new environments more efficiently. By learning how to learn, these techniques can rapidly incorporate necessary modifications and adapt to new scenarios with significantly reduced data requirements, making them well-suited for applications where data is scarce, or the environment is dynamically changing.

They can be broadly categorized into two types: *model-based* and *optimization-based* meta-learning. In the context of knowledge editing, model-based meta-learning focuses on learning a model that can be used to adapt to new data efficiently. This involves learning the structure of the model and its parameters such that it can be generalized to new data. In optimization-based meta-learning, the focus is on learning how to optimize the parameters of a model to adapt to new knowledge.

In literature, authors of "Editable Neural Networks" (ENN), [Sinitsin et al. \(2020\)](#), firstly exploit the meta-learning paradigm for knowledge editing to "learn to allow effective editing". The core idea behind "Editable Training" is to enforce the model parameters to be "prepared" for the editor function KE (Section 2.2), that in their experimentation is defined as Stochastic Gradient Descent (SGD) with up to k steps and learning rate α . In particular, they propose to train the starting model f_0 with a loss that at the same time encourages reliability and locality. That is obtained with two additional terms to the base loss, one that measures the success of an edit, cross-entropy, and another that assesses the distance in the output probability between the edited model and the original one, KL divergence. They prove to be successful on either computer vision datasets for classification [Krizhevsky et al. \(2009\)](#); [Deng et al. \(2009b\)](#) and Natural Adversarial Examples (NAE) [Hendrycks et al. \(2021\)](#) and machine translation [Cettolo et al. \(2014\)](#). They work with single non-successive and successive edits, with good performances on both types of edits. Nevertheless, as pointed out

by [Mitchell et al. \(2022a\)](#), the methodology requires to further training the base model before an edit with a pre-training step. That can be critical in scenarios with large models (LM) or when the memory available is a constraint.

Conversely, [De Cao et al. \(2021\)](#) they propose an optimization-based meta-learning, firstly employing a hypernetwork [Ha et al. \(2016\)](#) dubbed KnowledgeEditor with the objective to “learn to update” the parameters of another network. Therefore, since the knowledge editing task requires being locale except for the edit, they frame the learning task as a constrained optimization problem using a Bidirectional-LSTM [Schmidhuber et al. \(1997\)](#) and some downstream Feed-Forward Neural Networks (FFNN). Once trained, only the input edit feed the hypernetwork that predicts vectors to gate and shift the edit fine-tuning gradient of the starting network in respect to a certain weight matrix. Therefore, we can say that KnowledgeEditor learns how to modify the gradient in order to provide the property enumerated in Section 2.4. On the other hand, the training itself of the hypernetwork does not change the weights of the starting model as [Sinitisin et al. \(2020\)](#), but requires some original training samples to estimate the original output probability distribution. Indeed, as [Sinitisin et al. \(2020\)](#), the hypernetwork is trained with the sum of two loss components: the first one aims at providing reliability and generality with cross-entropy between semantically equivalent edits and predictions of the edited model and the second one provides locality with the KL divergence as [Sinitisin et al. \(2020\)](#). However, in addition to them, they propose to add a margin to KL and iteratively reduce it to progressively make the edit more local. They test their KnowledgeEditor on different NLP tasks, from fact-checking (FC) with a fine-tuned BERT model [Devlin et al. \(2018\)](#) with the FEVER dataset [Thorne et al. \(2018\)](#), to closed book question answering (QA) with BART model [Lewis et al. \(2020\)](#) on the Zero-Shot Relation Extraction (zsRE) dataset [Levy et al. \(2017\)](#). However, they only experimented with single non-successive changes. Nevertheless, authors of [Huang et al. \(2023\)](#) adopt KnowledgeEditor as a baseline and their experiments on FEVER and zsRE show that it fails to implement more than a couple of successive single edits. Indeed, as hypothesized by the authors, KnowledgeEditor hypernetwork is trained with the starting model f_0 and thus strongly coupled with the original parameters. As the editing proceeds, the model becomes more different from the initial one, resulting in their failure.

Building over [De Cao et al. \(2021\)](#), authors of [Mitchell et al. \(2022a\)](#) leverage on hypernetworks too in order to learn how to update weights of a starting model. However, while KnowledgeEditor trains a recurrent neural network to map the edit example into a rank-1 mask over the gradient, [Mitchell et al. \(2022a\)](#) hypernetwork, named MEND, directly maps the gradient into a new parameter update, retaining tractability by leveraging the low-rank form of the gradient. Indeed, the input of a MEND network is a decomposed gradient and the output is formed by some pseudoactivations and psudodelta that should encapsulate, reliability, generality, and locality. Indeed, as [De Cao et al. \(2021\)](#), they make use of cross-entropy between semantically equivalent edits and predictions of the edited model to enforce generability and reliability (edit loss), and KL divergence for locality without the margin (locality loss). In addition, they propose two further hyperparameters to perform a weighted sum of the two losses and to make learnable the edit learning rate coefficient. Each layer of the network itself consists of two consecutive blocks, both initialized to compute the exact identity function of the normalized decomposed gradient, using four matrices initialized with zero or Xavier uniform initialization [Glorot and Bengio \(2010\)](#). Finally, in order to edit multiple layers of the starting model with same matrices dimension, they propose to use the same MEND network, while applying a learned layer-specific scale and offset similar to [Perez et al. \(2018\)](#). As [De Cao et al. \(2021\)](#), they experiment their methodology only with NLP tasks using FEVER [Thorne et al. \(2018\)](#) and zsRE [Levy et al. \(2017\)](#) datasets. In addition, they also evaluate with GPT-style models, [Radford et al. \(2018\)](#), working with a custom version of Wikitext-103 [Merity et al. \(2016\)](#), named Wikitext Generation. They both experiment with non-sequential single or batch types of edits, showing large regression over 100 simultaneous edits. Finally, it is important to point out that as KnowledgeEditor, also MEND has the same limitations with successive edits, being strictly tied to the weights of the starting model. In fact, if pre-training is conducted on f_t , KE will exhibit significantly poorer performance with f_{t+1} . Moreover, as the weights of the edited model diverge from those of the pre-training, MEND will gradually deteriorate, ultimately losing all its editing capabilities.

Similar to MEND, MALMEN [Tan et al. \(2023\)](#) utilizes a hypernetwork to generate parameter shifts that can be applied to the language model. However, MALMEN distinguishes itself by formulating the parameter shift aggregation as a least squares problem, which it solves using the normal equation. This allows MALMEN to effectively combine the parameter shifts corresponding to different facts, mitigating the potential cancellation effects that can arise when simply summing the shifts. Additionally, MALMEN separates the computation between the hypernetwork and the language model, enabling the use of arbitrary batch sizes on both components. This memory-economic training strategy permits MALMEN to scale to editing thousands of facts simultaneously, substantially outperforming the editing capabilities of MEND. The paper provides thorough empirical evaluations demonstrating MALMEN’s superior scalability and effectiveness across various language model architectures and knowledge-intensive NLP tasks, including closed book fact verification on the FEVER dataset [Thorne et al. \(2018\)](#) and question answering on the zsRE dataset [Levy et al. \(2017\)](#) for BERT-base, GPT-2, T5-XL (2.8B), and GPT-J (6B). Finally, it is important to point out that as KnowledgeEditor, and MEND, also MALMEN has the same limitations with successive edits, being strictly tied to the weights of the starting model.




Figure 2: Scaling curves showing three different evaluation metrics with an increased number of non-successive batch edits for three different KE methodologies: MEND, ROME, and MEMIT. Results are computed using COUNTERFACT and GPT-J. Locality is shown not as drawdown (DD), but as its complementary specificity over a neighborhood of samples Meng et al. (2022b). ROME and MEND performs well up to ten edits, but rapidly degrade, losing almost all SR before batches of 1k. On the other hand, MEMIT performs well with considerable large batches of edits. Adapted from Meng et al. (2022a).

4.3 Direct Model Editing

This family of approaches, known as direct model editing, aims to directly edit the weights or parameters within the target neural network, enabling effective implementation of a predefined set of changes while minimizing computational overhead. These methodologies allow injecting knowledge by modifying only certain parameters of the model, making them particularly suitable for foundational models Chang et al. (2023). By directly manipulating the variables of a model, most direct model editing techniques build upon efforts to localize and understand the internal mechanisms within models, striving to attribute knowledge acquired during training to specific neurons or parameters in the network Elhage et al. (2021); Dar et al. (2023); Mitchell et al. (2022b). Consequently, these approaches endeavor to edit the activations or values of identified neurons or parameters to reflect the desired changes.

However, it is crucial to recognize that the applicability and effectiveness of direct model editing techniques may vary depending on the underlying neural network architecture and the nature of the knowledge or changes to be incorporated. While extensive research has been conducted on specific architectures such as Multilayer Perceptrons (MLPs) and Transformers, the generalization of these techniques to alternative architectures or problem domains warrants further investigation. Therefore, an in-depth analysis and exploration of the applicable scenarios for direct model editing techniques are beneficial, considering factors such as the model architecture, the type of knowledge or changes to be incorporated, and the computational constraints of the target environment.

Geva, Mor, et al. Tay et al. (2022), firstly identify the MLP layers in a masked LM transformer as key-value memories of entities and information associated with that entity, Sukhbaatar et al. (2015). Building on this finding, Dai et al. (2022) demonstrate a method to edit facts in BERT models Devlin et al. (2018); they propose the concept of knowledge neurons to study how factual knowledge is stored in pretrained Transformers. Specifically, the authors examine the fill-in-the-blank cloze task for BERT and propose a knowledge attribution method to identify the neurons that express a given relational fact. They find that the activation of such knowledge neurons is positively correlated to the expression of their corresponding facts. Therefore, they present a preliminary methodology that leverages knowledge neurons to edit factual knowledge in Transformers, even without any fine-tuning. The authors perform a knowledge surgery for pretrained Transformers by directly modifying the corresponding parameters in feed-forward networks. Such surgery shows promising results, keeping a moderate influence on other knowledge. The methodology proposed can be used to perform both single and multiple edits at the same time. However, the authors only experimented with single edits in their paper, focusing on specifically on factual knowledge.

Factual knowledge is also the focus of Meng Kevin, et al. work, Meng et al. (2022b), but exclusively experimenting on autoregressive language modelling, such as GPT-style models, Radford et al. (2018). The study explores the storage and direct manipulation of factual associations in transformer models. These associations are modeled as tuples, represented as $t = (s, r, o)$, comprising a subject s , a relation r , and an object o , which connect to form a factual relationship. The research focuses into understanding how these tuples are stored within transformer models and investigates methods to facilitate their direct editing. Firstly, they trace the causal effects of hidden state activations within GPT using causal mediation analysis Pearl (2001) as previously done by other works Vig et al. (2020) to identify the specific modules that mediate recall of a fact about a subject. Their analysis reveals that feedforward MLPs at a range of middle

layers are decisive when processing the last token of the subject name. Secondly, they test their previous findings by introducing a novel direct model editing methodology: Rank-One Model Editing method (ROME). The algorithm aims at producing single, not successive edits by altering the parameters that determine a feedforward layer’s behavior. As stated by the authors, despite the simplicity of the intervention, ROME is similarly effective to other model-editing approaches achieving good generalization and locality simultaneously, whereas previous approaches sacrifice one or the other. Furthermore, they introduce a new dataset, CounterFact (derived from a set of true facts from WikiData), for evaluating counterfactual edits in language models. It contains 21,919 records with a diverse set of subjects, relations, and linguistic variations. Finally, they propose to monitor a further metrics, *fluency*, to evaluate text generation’s repetitiveness by measuring the weighted average of bi- and tri-gram entropies [Zhang et al. \(2018\)](#).

ROME results are comparable with other knowledge editing meta-learning techniques. Nevertheless, the methodology is limited to work with single non-successive edits. To address this problem, the same authors modified the original algorithm to overcome this problem [Meng et al. \(2022a\)](#). The method, called MEMIT, is able to scale up to thousands of associations for GPT-J (6B) and GPT-NeoX (20B), exceeding prior work by an order of magnitude. As ROME, MEMIT works by first identifying the critical layers in the LLM that are responsible for storing the memories that need to be updated. Then, it uses a rank-one modification of the MLP weights of these layers to directly write the new memories into the model. This process is repeated for each memory that needs to be updated. The authors evaluated MEMIT on a variety of tasks, including factual question answering, natural language inference, and summarization. They found that MEMIT was able to significantly improve the performance of the LLM on these tasks, even when the number of memories that were updated was large. In addition to its scalability, MEMIT also has several other advantages over previous methods for updating LLMs. First, it is relatively efficient, requiring only a few minutes to update a large model with thousands of memories. Second, it is very accurate, with the updated models achieving near-perfect accuracy on the tasks that were evaluated. Third, MEMIT is generalizable, with the updated models able to perform well on a variety of tasks. [Figure 2](#) compares ROME, MEMIT, and one of the meta-learning techniques described in [section 4.2](#). It is possible to appreciate the improvement brought with MEMIT and the clear collapse of the MEND HyperNetwork-based editor with large batches. Overall, MEMIT is a valuable tool for improving the performance of these models. However, it is important to point out that, even if there are works that try to expand the scope of the methodology [Gupta et al. \(2023\)](#), MEMIT is limited to modifying factual association and that there is not a clear path to scale it to different knowledge types. The same applies for PMET [Li et al. \(2024\)](#) which goes a step further by also simultaneously optimizing the hidden states of the Multi-Head Self-Attention (MHSA) component. This is based on the insight that MHSA encodes certain general knowledge extraction patterns that can be leveraged to enable more precise editing of the model. On the other hand, it is worth noticing that recent research suggests that casual tracing or other knowledge localization methodologies for identifying which parameters to update are surprisingly unreliable also for factual knowledge [Hase et al. \(2023\)](#). The paper argues that locating the source of a fact does not necessarily translate to effective editing, as modifying that specific location may not change the model’s output. This surprising finding indicates our understanding of how knowledge is stored in complex models is incomplete.

4.4 Architectural Strategies

Architecture strategies represent a distinct family of methodologies that diverge from approaches that directly or indirectly modify the weights of the target model. Instead, these strategies concentrate on manipulating or augmenting the original architecture to patch and modify the network’s existing knowledge. This approach can be advantageous in scenarios where the set of original weights is inaccessible or when it is deemed unsafe to directly manipulate the main model. This indirect manipulation can be particularly useful when dealing with highly sensitive or proprietary models, where direct weight modification may be prohibited or could potentially compromise the model’s integrity. Additionally, architecture strategies can offer a more interpretable and controllable means of knowledge adaptation, as the introduced architectural changes can often be mapped to specific functional aspects of the model. However, it is always crucial to consider the trade-offs between the potential benefits of architectural manipulation and the computational overhead associated with it in terms of latency.

A first editor function of this type was proposed by [Sinitin et al. \(2020\)](#), that taking inspiration from Conditional Neural Processes (CNP) [Garnelo et al. \(2018\)](#), they propose a specialized architecture that performs edits by adding a special condition vector to intermediate activations. The vector is generated by an additional encoder layer and provides information to effectively steer the final prediction of the network. This approach is very similar to more generic memory-augmented models that introduce memory mechanisms to enhance neural networks [Graves et al. \(2014\)](#); [Santoro et al. \(2016\)](#). On the same line of experimentation, [Mitchell et al. \(2022b\)](#) proposed an approach, called SERAC (Semi-Parametric Editing with a Retrieval-Augmented Counterfactual Model), which stores edits in an explicit memory and learns to reason over them to modulate the base model’s predictions as needed. This allows SERAC to be more expressive than other model editors, and to produce more accurate predictions for test inputs that are not closely related to the edit. That is achieved using three components: a memory, a classifier,

and a counterfactual model. Users adds edit in the memory and the classifier decides whether the memory contains inputs relevant to processing them. If the classifier determines that a relevant edit example exists, the input and edit example are passed to the counterfactual model that is responsible for making the prediction. They evaluate SERAC on numerous tasks, focusing exclusively on large language models with single and batch non-successive edits.

On the same line of SERAC, authors of [Dong et al. \(2022\)](#) presented a memory-based methodology called CaliNet. It is a method for calibrating specifically factual knowledge in pretrained language models (PLM) without re-training from scratch. CaliNet first detects whether a PLM has learned a fact correctly by comparing its scores for the correct fact and a set of distractor facts. If the PLM has not learned the fact correctly, CaliNet then uses a lightweight method to add and adapt new parameters to the PLM for that specific fact. CaliNet has two main components: first, a contrastive knowledge assessment (CKA) module that detects whether a PLM has learned a fact correctly. The CKA module works by comparing the PLM’s scores for the correct fact and a set of distractor facts. If the PLM assigns a higher score to the correct fact than to the distractor facts, then the CKA module concludes that the PLM has learned the fact correctly. Secondly, a factual knowledge calibration (FKC) module adds and adapts new parameters to the PLM for a specific fact. The FKC module works by first creating a new set of parameters for the PLM that are specific to the given fact. These new parameters are then initialized with the PLM’s original parameters for the fact. The FKC module then fine-tunes the new parameters on a dataset of factual examples that include the given fact. Using a custom dataset based on ParaRel set [Elazar et al. \(2021\)](#), CaliNet has shown to be effective at calibrating factual knowledge in PLMs. In experiments on the knowledge probing task, CaliNet was able to improve the factual accuracy of PLMs by up to 20%.

All previous works have tackled only non-successive edits, making only one edit or batch of edits at a time. However, as pointed out by [Huang et al. \(2023\)](#), the one-mistake-fixing scenario is not an accurate abstraction of the real-world challenge. Therefore, they extend the scenario into successive edits, introducing a novel model editor, Transformer-Patcher (T-Patcher), that can shift the behavior of transformer-based models by simply adding and training a few neurons in the last Feed-Forward Network layer. Being at the end of the model, the new neurons have access to the output of all the previous layers. This allows the new neurons to learn how to correct mistakes that are made in earlier layers of the model. Overall, the proposed methodology allows for targeted shifts in the behavior of the model, akin to other fine-tuned oriented methodologies like LoRA [Hu et al. \(2022\)](#). Transformer-Patcher is able to successively correct up to thousands of errors while generalizing to their equivalent inputs and retaining the model’s accuracy on irrelevant inputs. This is in contrast to previous methods, which either fail to make a sequence of edits or to remember previous edits. The work evaluates Transformer-Patcher on both classification and generation tasks, showing that Transformer-Patcher can achieve state-of-the-art performance for single successive edits. However, despite the expansive scope of the methodology, their approach and implementation are highly architecture-specific, relying heavily also on large sources of data and unrelated inputs.

To this end, Hartvigsen, Thomas, et al. [Hartvigsen et al. \(2022\)](#), proposes GRACE, a methodology that requires only samples of the edits and that can perform successive edits, ensuring minimal impact on unrelated inputs. GRACE works by caching a chosen layer’s activations in an adaptive codebook as edits stream in. When a new edit is received, GRACE retrieves the corresponding activations from the codebook and uses them to update the model’s predictions. This allows GRACE to make targeted edits to the model’s behavior without affecting its overall performance. The authors evaluated GRACE on a variety of tasks, including text classification, natural language inference, and question answering. They showed that GRACE was able to significantly improve the performance of LLMs on streaming errors, while minimally affecting their performance on unrelated inputs. Finally, the authors do not explicitly mention the impact of GRACE on latency. However, it is possible that GRACE could have a small impact on latency, as it requires caching activations in an adaptive codebook.

5 Conclusion and Future Directions

In this survey, we organize recent progress of the increasingly growing field of knowledge editing. We began formalizing it under a common definition that seamlessly connects the slightly different facets of each work presented in the literature so far. Indeed, being a very recent branch of research, each author tried to differently bend definitions to better accommodate their methodologies. So, following the first definitions given by [Sinitsin et al. \(2020\)](#) and the more recent adaptation by [Huang et al. \(2023\)](#), we formally define knowledge editing as the task of modifying the knowledge of a target model in a non-sequential or sequential manner, without significantly affecting its original performance. This makes knowledge editing closely related to the much more well-known branch of continuous learning, or lifelong learning, as well as the emerging fields of machine unlearning [Bourtole et al. \(2021\)](#); [Nguyen et al. \(2022\)](#); [Si et al. \(2023\)](#) and parameter-efficient fine-tuning [Fu et al. \(2023\)](#); [Hu et al. \(2022\)](#); [Han et al. \(2024\)](#); [Yu et al. \(2024\)](#), including most notable variants like model editing via task arithmetic [Ilharco et al. \(2022\)](#); [Hendel et al. \(2023\)](#). In the following section, we briefly discuss the intersections and connections between knowledge editing and related disciplines.

5.1 Knowledge editing and related fields of research

Machine unlearning researches focus on removing specific data samples or knowledge from a pre-trained model without retraining the entire model from scratch. This task is gaining relevance due to its applications in data privacy [Chen et al. \(2021\)](#), security [Cao and Yang \(2015\)](#), and compliance with regulations such as GDPR [Sai et al. \(2023\)](#). However, while knowledge editing primarily focuses on modifying or adding knowledge to a model, machine unlearning can be seen as a complementary task that aims to selectively remove knowledge from a model. Both tasks share the common goal of avoiding complete retraining while maintaining the model’s predictive capability, but tracking different metrics and evaluation benchmarks.

Knowledge editing and machine unlearning can be viewed as specialized cases of continual learning [Mundt et al. \(2023\)](#), focusing on applying targeted, often non-uniformly distributed edits to a network’s existing knowledge [Henn et al. \(2021\)](#). In contrast, continual learning encompasses a broader scope, seeking general methodologies to incrementally expand a network’s knowledge base, enabling it to perform an increasing array of tasks and skills over time. The key distinction lies in their objectives: knowledge editing aims for precise modifications to specific knowledge elements, while continual learning emphasizes the gradual accumulation of knowledge and skills without compromising previously learned information. Despite these differences, both fields share significant challenges, particularly in modifying neural networks without disrupting existing capabilities. Both contend with issues such as catastrophic forgetting [Ratcliff \(1990\)](#), where new learning can overwrite previously acquired knowledge. However, the approaches to addressing these challenges often diverge. Knowledge editing techniques typically employ localized, precise modifications, while continual learning methods often utilize more global strategies to balance retention of old knowledge with acquisition of new information. The intersection of these fields is particularly evident in scenarios where techniques from one inspire developments in the other. For example, some knowledge editing approaches adapt continual learning strategies for mitigating catastrophic forgetting (Section 4.1), tailoring them for more targeted edits. Conversely, insights from knowledge editing about precise neural network modifications can inform the development of more fine-grained continual learning algorithms.

On the other hand, parameter-efficient fine-tuning (PET) techniques are another related field that involves modifying pre-trained models to create new models with desired capabilities. PET aims to enable efficient adaptation by updating only a minimal subset of model parameters, rather than fine-tuning all parameters. Notable PET techniques include addition-based methods like adapters [Houlsby et al. \(2019\)](#), specification-based methods like LoRA [Hu et al. \(2022\)](#), and reparameterization-based methods. A interesting variant of PET is task arithmetic [Ilharco et al. \(2022\)](#), which involves operations such as model addition, subtraction, multiplication, and permutation to modify or reshape the knowledge encoded in pre-trained models. By manipulating entire model representations in this arithmetic fashion, task arithmetic provides a way to create tailored models for specific tasks or domains without retraining from scratch, similar to knowledge editing techniques. However, task arithmetic and PET more in general are typically applied to enhance task performance rather than edit knowledge specifically. The efficacy of existing PET methods for knowledge editing remains largely unexplored. Indeed, while task arithmetic and other PET techniques hold promise for efficient model adaptation, they differ from knowledge editing in their primary focus on improving downstream task metrics rather than directly modifying a model’s acquired knowledge. Investigating how to leverage PET methods for precise and targeted knowledge updates presents an interesting direction for future work in the knowledge editing space.

5.2 Future Directions and Risks

Building on the more formal definition, we presented a distilled summary of the most relevant works in the literature at the time of writing. We proposed to categorize these works into four families: regularization techniques, meta-learning, direct model editing, and architectural strategies. We discussed each family in turn, highlighting their intrinsic characteristics and limitations. We also summarized the most frequent field of application, tasks, and datasets that have been tackled in each family, for quick reference. We did not specifically deep dive into the works where knowledge editing could emerge as an additional benefit of the proposed methodologies [Hewitt et al. \(2023\)](#) but it is worth noting that future expansions of similar survey research could encompass these aspects for a more comprehensive analysis.

Overall, we have presented a rapidly expanding field of research driven by the current trend of foundational models [Zhou et al. \(2023\)](#). The advancements in this area have led to a significant increase in the development of tools and methodologies to effectively harness the intrinsic knowledge of these models. As we move forward, knowledge editing is poised to become a critical factor in leveraging the power of these models for various industrial applications. However, several key challenges and future directions remain to be addressed:

- Improving the efficiency and scalability of knowledge editing techniques to handle large-scale foundational models with constrained computational resources. This challenge is particularly pressing as models continue

to grow in size and complexity [Dubey et al. \(2024\)](#). Future research could explore techniques such as sparse editing or hierarchical knowledge representations to enable more efficient updates.

- Enhancing methodologies to introspect the different types of distributed knowledge stored inside neural networks, to enable better understanding and control over the changes made to models [Templeton \(2024\)](#). This could involve developing novel visualization techniques or interpretability methods specifically tailored for knowledge editing tasks.
- Developing robust benchmarks specifically tailored for knowledge cross-domains editing tasks, to better assess the performance and impact of different techniques. These benchmarks should not only measure the effectiveness of edits but also their impact on the model’s overall performance and generalization capabilities.
- Investigating the potential of knowledge editing for improving the robustness, fairness, and ethical behavior of AI models, by correcting biases, removing harmful tendencies, or incorporating desired values and principles. This direction could lead to the development of "ethical editing" frameworks that ensure AI systems align with human values and societal norms.
- Developing knowledge editing techniques to incorporate edits at runtime, allowing to create systems able to quickly adapt to errors and changes. This could pave the way for more adaptive and flexible AI systems that can learn and update their knowledge on-the-fly.
- Investigating the potential of knowledge editing in multi-modal models [Baltrušaitis et al. \(2018\)](#), where information is represented across different modalities (e.g., text, images, audio). This research direction could lead to more versatile and comprehensive knowledge editing techniques applicable to a wider range of AI systems.

Finally, it is vital to acknowledge that the power of knowledge editing also brings inherent risks that must not be overlooked. While editing models can correct their behavior and improve their utility, it can also be exploited for harmful purposes. In particular, sophisticated editing algorithms may enable malicious actors to deliberately incorporate backdoors, vulnerabilities, hidden behaviors, or harmful tendencies into the models. This concern becomes even more critical for methodologies that edit weights without providing sufficient interpretability of the applied changes. This dual use is a common risk for many machine learning technologies, and only proactive efforts to develop robust security measures and ethical guidelines can help to mitigate these potential risks. Future research should focus on developing "edit verification" techniques that can detect and prevent malicious edits, as well as establishing standardized protocols for auditing and certifying edited models.

In conclusion, knowledge editing in AI has emerged as a critical field, offering transformative potential for enhancing AI capabilities while simultaneously raising significant challenges and ethical considerations. As researchers and practitioners, we must strive to balance the pursuit of technological advancements with a strong awareness of their broader societal impacts, ensuring that the evolution of knowledge editing techniques contributes positively to the development of safe, reliable, and beneficial AI systems.

References

- Joseph Cichon and Wen-Biao Gan. Branch-specific dendritic ca^{2+} spikes cause persistent synaptic plasticity. *Nature*, 520(7546):180–185, 2015.
- Roger Ratcliff. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285, 1990.
- Marcus K Benna and Stefano Fusi. Computational principles of synaptic memory consolidation. *Nature neuroscience*, 19(12):1697–1706, 2016.
- Guang Yang, Feng Pan, and Wen-Biao Gan. Stably maintained dendritic spines are associated with lifelong memories. *Nature*, 462(7275):920–924, 2009.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. Transformer-patcher: One mistake worth one neuron. *arXiv preprint arXiv:2301.09785*, 2023.
- Elena Voita, Rico Sennrich, and Ivan Titov. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives. In *2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, pages 4387–4397. Association for Computational Linguistics (ACL), 2019.

- Yonatan Belinkov and James Glass. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72, 2019.
- Anton Sinitin, Vsevolod Plokhotnyuk, Dmitriy Pyrkun, Sergei Popov, and Artem Babenko. Editable neural networks. *arXiv preprint arXiv:2004.00345*, 2020.
- Justyna Sarzynska-Wawer, Aleksander Wawer, Aleksandra Pawlak, Julia Szymanowska, Izabela Stefaniak, Michal Jarkiewicz, and Lukasz Okruszek. Detecting formal thought disorder by deep contextualized word representations. *Psychiatry Research*, 304:114135, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyers, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. Mind the gap: Assessing temporal generalization in neural language models. *Advances in Neural Information Processing Systems*, 34:29348–29363, 2021.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. Editing factual knowledge in language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, 2021.
- Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix Yu, and Sanjiv Kumar. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*, 2020.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model editing at scale. In *International Conference on Learning Representations*, 2022a. URL <https://openreview.net/pdf?id=ODcZxeWfOPt>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Saleh Soltan, Shankar Ananthkrishnan, Jack FitzGerald, Rahul Gupta, Wael Hamza, Haidar Khan, Charith Peris, Stephen Rawls, Andy Rosenbaum, Anna Rumshisky, et al. Alexatm 20b: Few-shot learning using a large-scale multilingual seq2seq model. *arXiv preprint arXiv:2208.01448*, 2022.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. Zero-shot relation extraction via reading comprehension. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, 2017.
- Siyuan Cheng, Ningyu Zhang, Bozhong Tian, Xi Chen, Qingbin Liu, and Huajun Chen. Editing language model-based knowledge graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17835–17843, 2024.
- J Thorne, A Vlachos, C Christodoulopoulos, and A Mittal. Fever: a large-scale dataset for fact extraction and verification. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Sheffield, 2018.
- Mohammad Shehab, Laith Abualigah, Qusai Shambour, Muhannad A Abu-Hashem, Mohd Khaled Yousef Shambour, Ahmed Izzat Alsalibi, and Amir H Gandomi. Machine learning in medical applications: A review of state-of-the-art methods. *Computers in Biology and Medicine*, 145:105458, 2022.
- Mohsen Soori, Behrooz Arezoo, and Roza Dastres. Artificial intelligence, machine learning and deep learning in advanced robotics, a review. *Cognitive Robotics*, 2023.
- Abhinav Sharma, Arpit Jain, Prateek Gupta, and Vinay Chowdary. Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, 9:4843–4873, 2020.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities. *arXiv preprint arXiv:2305.13172*, 2023.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*, 2022a.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. Memory-based model editing at scale. In *International Conference on Machine Learning*, pages 15817–15831. PMLR, 2022b.
- Thomas Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. Aging with grace: Lifelong model editing with discrete key-value adaptors. *arXiv preprint arXiv:2211.11031*, 2022.
- Tom Michael Mitchell et al. *Machine learning*, volume 1. McGraw-hill New York, 2007.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022b.
- Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International conference on machine learning*, pages 1704–1713. PMLR, 2018.
- Richard Szeliski. *Computer vision: algorithms and applications*. Springer Nature, 2022.
- Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, Eftychios Protopapadakis, et al. Deep learning for computer vision: A brief review. *Computational intelligence and neuroscience*, 2018, 2018.
- Niall O’Mahony, Sean Campbell, Anderson Carvalho, Suman Harapanahalli, Gustavo Velasco Hernandez, Lenka Krpalkova, Daniel Riordan, and Joseph Walsh. Deep learning vs. traditional computer vision. In *Advances in Computer Vision: Proceedings of the 2019 Computer Vision Conference (CVC), Volume 1 1*, pages 128–144. Springer, 2020.
- Michael Firman. Rgb-d datasets: Past, present and future. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 19–31, 2016.
- Matthew Sotoudeh and Aditya V Thakur. Provable repair of deep neural networks. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, pages 588–603, 2021.
- Norman Mu and Justin Gilmer. Mnist-c: A robustness benchmark for computer vision. *arXiv preprint arXiv:1906.02337*, 2019.
- Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. URL <https://api.semanticscholar.org/CorpusID:18268744>.
- 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, 2009a. doi:10.1109/CVPR.2009.5206848.
- Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 312–321, 2019.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi:10.1007/s11263-015-0816-y.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Jacob Eisenstein. *Introduction to Natural Language Processing*. MIT Press, 2019.
- Daniel W Otter, Julian R Medina, and Jugal K Kalita. A survey of the usages of deep learning for natural language processing. *IEEE transactions on neural networks and learning systems*, 32(2):604–624, 2020.
- Jack FitzGerald, Shankar Ananthkrishnan, Konstantine Arkoudas, Davide Bernardi, Abhishek Bhagia, Claudio Delli Bovi, Jin Cao, Rakesh Chada, Amit Chauhan, Luoxin Chen, Anurag Dwarakanath, Satyam Dwivedi, Turan Goyayev, Karthik Gopalakrishnan, Thomas Gueudre, Dilek Hakkani-Tur, Wael Hamza, Jonathan J. Hüser, Kevin Martin Jose, Haidar Khan, Beiye Liu, Jianhua Lu, Alessandro Manzotti, Pradeep Natarajan, Karolina Owczarzak, Gokmen

- Oz, Enrico Palumbo, Charith Peris, Chandana Satya Prakash, Stephen Rawls, Andy Rosenbaum, Anjali Shenoy, Saleh Soltan, Mukund Harakere Sridhar, Lizhen Tan, Fabian Triefenbach, Pan Wei, Haiyang Yu, Shuai Zheng, Gokhan Tur, and Prem Natarajan. Alexa teacher model: Pretraining and distilling multi-billion-parameter encoders for natural language understanding systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, page 2893–2902, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393850. doi:[10.1145/3534678.3539173](https://doi.org/10.1145/3534678.3539173). URL <https://doi.org/10.1145/3534678.3539173>.
- Haifeng Wang, Jiwei Li, Hua Wu, Eduard Hovy, and Yu Sun. Pre-trained language models and their applications. *Engineering*, 2022.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206, 2022.
- Pedro A Ortega, Markus Kunesch, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Joel Veness, Jonas Buchli, Jonas Degraeve, Bilal Piot, Julien Perolat, et al. Shaking the foundations: delusions in sequence models for interaction and control. *arXiv preprint arXiv:2110.10819*, 2021.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023.
- Tal Schuster, Adam Fisch, and Regina Barzilay. Get your vitamin C! robust fact verification with contrastive evidence. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online, June 2021. Association for Computational Linguistics. doi:[10.18653/v1/2021.naacl-main.52](https://doi.org/10.18653/v1/2021.naacl-main.52). URL <https://aclanthology.org/2021.naacl-main.52>.
- Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*, 2020.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. T-REX: A large scale alignment of natural language with knowledge base triples. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA). URL <https://www.aclweb.org/anthology/L18-1544>.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019. doi:[10.1162/tacl_a_00276](https://doi.org/10.1162/tacl_a_00276). URL <https://aclanthology.org/Q19-1026>.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. FairLex: A multilingual benchmark for evaluating fairness in legal text processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4389–4406, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi:[10.18653/v1/2022.acl-long.301](https://doi.org/10.18653/v1/2022.acl-long.301). URL <https://aclanthology.org/2022.acl-long.301>.
- John C Knight. Safety critical systems: challenges and directions. In *Proceedings of the 24th international conference on software engineering*, pages 547–550, 2002.
- Abhishek Gupta, Alagan Anpalagan, Ling Guan, and Ahmed Shaharyar Khwaja. Deep learning for object detection and scene perception in self-driving cars: Survey, challenges, and open issues. *Array*, 10:100057, 2021.
- Ahmet İlker Tekkeşin et al. Artificial intelligence in healthcare: past, present and future. *Anatol J Cardiol*, 22(Suppl 2): 8–9, 2019.
- Banavar Sridhar. Applications of machine learning techniques to aviation operations: Promises and challenges. In *2020 International Conference on Artificial Intelligence and Data Analytics for Air Transportation (AIDA-AT)*, pages 1–12. IEEE, 2020.
- Kyle D. Julian, Jessica Lopez, Jeffrey S. Brush, Michael P. Owen, and Mykel J. Kochenderfer. Policy compression for aircraft collision avoidance systems. In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, pages 1–10, 2016. doi:[10.1109/DASC.2016.7778091](https://doi.org/10.1109/DASC.2016.7778091).
- Kyle D Julian and Mykel J Kochenderfer. Guaranteeing safety for neural network-based aircraft collision avoidance systems. In *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*, pages 1–10. IEEE, 2019.
- Feisi Fu and Wenchao Li. Sound and complete neural network repair with minimality and locality guarantees. *arXiv preprint arXiv:2110.07682*, 2021.

- Zhen Liang, Taoran Wu, Changyuan Zhao, Wanwei Liu, Bai Xue, Wenjing Yang, and Ji Wang. Repairing deep neural networks based on behavior imitation. *arXiv preprint arXiv:2305.03365*, 2023.
- Mykel J Kochenderfer and JP Chryssanthacopoulos. Robust airborne collision avoidance through dynamic programming. *Massachusetts Institute of Technology, Lincoln Laboratory, Project Report ATC-371*, 130, 2011.
- Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Formal security analysis of neural networks using symbolic intervals. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1599–1614, 2018.
- Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *Computer Aided Verification: 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I 30*, pages 97–117. Springer, 2017.
- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
- Zirui Liu, Zhimeng Jiang, Shaochen Zhong, Kaixiong Zhou, Li Li, Rui Chen, Soo-Hyun Choi, and Xia Hu. Editable graph neural network for node classifications. *arXiv preprint arXiv:2305.15529*, 2023.
- Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3:127–163, 2000.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018.
- Martin Mundt, Yongwon Hong, Iuliia Pliushch, and Visvanathan Ramesh. A wholistic view of continual learning with deep neural networks: Forgotten lessons and the bridge to active and open world learning. *Neural Networks*, 2023.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- Mengye Ren, Eleni Triantafyllou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*, 2018.
- Jiatao Gu, Yong Wang, Yun Chen, Kyunghyun Cho, and Victor OK Li. Meta-learning for low-resource neural machine translation. In *2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 3622–3631. Association for Computational Linguistics, 2020.
- Rein Houthoofd, Yuhua Chen, Phillip Isola, Bradly Stadie, Filip Wolski, OpenAI Jonathan Ho, and Pieter Abbeel. Evolved policy gradients. *Advances in Neural Information Processing Systems*, 31, 2018.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 2009. *URL* <http://www.cs.toronto.edu/kriz/cifar.html>, 5, 2009.
- 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. Ieee, 2009b.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report on the 11th iwslt evaluation campaign. In *Proceedings of the 11th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 2–17, 2014.
- David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Jürgen Schmidhuber, Sepp Hochreiter, et al. Long short-term memory. *Neural Comput*, 9(8):1735–1780, 1997.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.

- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual Reasoning with a General Conditioning Layer. In *AAAI Conference on Artificial Intelligence*, New Orleans, United States, February 2018. URL <https://inria.hal.science/hal-01648685>.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Chenmien Tan, Ge Zhang, and Jie Fu. Massive editing for large language models via meta learning. *arXiv preprint arXiv:2311.04661*, 2023.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 2023.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1, 2021.
- Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. Analyzing transformers in embedding space. In *Annual Meeting of the Association for Computational Linguistics*, 2023.
- Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 35: 21831–21843, 2022.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. End-to-end memory networks. *Advances in neural information processing systems*, 28, 2015.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, 2022.
- Judea Pearl. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, UAI’01, page 411–420, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558608001.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. Investigating gender bias in language models using causal mediation analysis. *Advances in neural information processing systems*, 33:12388–12401, 2020.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. Generating informative and diverse conversational responses via adversarial information maximization. *Advances in Neural Information Processing Systems*, 31, 2018.
- Anshita Gupta, Debanjan Mondal, Akshay Krishna Sheshadri, Wenlong Zhao, Xiang Lorraine Li, Sarah Wiegrefe, and Niket Tandon. Editing commonsense knowledge in gpt. *arXiv preprint arXiv:2305.14956*, 2023.
- Xiaopeng Li, Shasha Li, Shezheng Song, Jing Yang, Jun Ma, and Jie Yu. Pmet: Precise model editing in a transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18564–18572, 2024.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models. *arXiv preprint arXiv:2301.04213*, 2023.
- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850. PMLR, 2016.
- Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui, and Lei Li. Calibrating factual knowledge in pretrained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5937–5947, 2022.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. Measuring and improving consistency in pretrained language models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031, 2021.

- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 141–159. IEEE, 2021.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning. *arXiv preprint arXiv:2209.02299*, 2022.
- Nianwen Si, Hao Zhang, Heyu Chang, Wenlin Zhang, Dan Qu, and Weiqiang Zhang. Knowledge unlearning for llms: Tasks, methods, and challenges. *arXiv preprint arXiv:2311.15766*, 2023.
- Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 12799–12807, 2023.
- Zeyu Han, Chao Gao, Jinyang Liu, Sai Qian Zhang, et al. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024.
- Lang Yu, Qin Chen, Jie Zhou, and Liang He. Melo: Enhancing model editing with neuron-indexed dynamic lora. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19449–19457, 2024.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*, 2022.
- Roe Hendel, Mor Geva, and Amir Globerson. In-context learning creates task vectors. *arXiv preprint arXiv:2310.15916*, 2023.
- Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang. When machine unlearning jeopardizes privacy. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, pages 896–911, 2021.
- Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pages 463–480. IEEE, 2015.
- Siva Sai, Uday Mittal, Vinay Chamola, Kaizhu Huang, Indro Spinelli, Simone Scardapane, Zhiyuan Tan, and Amir Hussain. Machine un-learning: An overview of techniques, applications, and future directions. *Cognitive Computation*, pages 1–25, 2023.
- Thomas Henn, Yasukazu Sakamoto, Clément Jacquet, Shunsuke Yoshizawa, Masamichi Andou, Stephen Tchen, Ryosuke Saga, Hiroyuki Ishihara, Katsuhiko Shimizu, Yingzhen Li, et al. A principled approach to failure analysis and model repairment: Demonstration in medical imaging. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24*, pages 509–518. Springer, 2021.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR, 2019.
- John Hewitt, John Thickstun, Christopher D Manning, and Percy Liang. Backpack language models. *arXiv preprint arXiv:2305.16765*, 2023.
- Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *arXiv preprint arXiv:2302.09419*, 2023.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Adly Templeton. *Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet*. Anthropic, 2024.
- 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.