

Transformer in Touch: A Survey

Jing Gao¹, Ning Cheng¹, Bin Fang², and Wenjuan Han^{1,3,4}✉

¹ Beijing Key Lab of Traffic Data Analysis and Mining, Beijing Jiaotong University, Beijing, China

² School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing, China

³ China Railway Design Corporation, Tianjin, China

⁴ National Engineering Research Center for Digital Construction and Evaluation of Urban Rail-Transit, Tianjin, China
wjhan@bjtu.edu.cn

Abstract. The Transformer model, initially achieving significant success in the field of natural language processing, has recently shown great potential in the application of tactile perception. This review aims to comprehensively outline the application and development of Transformers in tactile technology. We first introduce the two fundamental concepts behind the success of the Transformer: the self-attention mechanism and large-scale pre-training. Then, we delve into the application of Transformers in various tactile tasks, including but not limited to object recognition, cross-modal generation, and object manipulation, offering a concise summary of the core methodologies, performance benchmarks, and design highlights. Finally, we suggest potential areas for further research and future work, aiming to generate more interest within the community, tackle existing challenges, and encourage the use of Transformer models in the tactile field.

Keywords: Tactile, Self-attention, Transformers, Self-supervision

1 Introduction

In recent years, the Transformer model [1] has caused a sensation in the world of deep learning. In particular, the Transformer model has revolutionized diverse fields with its attention mechanism. This mechanism allows the model to focus on different parts of the input sequence, capturing long-range dependencies and improving performance on various tasks. The Transformer's ability to handle sequential data has paved the way for significant advancements, especially in the fields of natural language processing (NLP) and computer vision (CV) [2].

In the field of NLP, Transformer continues to be a major driving force for technological advancement [3, 4]. Apart from the well-known BERT [5] and GPT-3 [6], newer Transformer-based models have shown even stronger performance. For instance, GPT-4 [7] offers more refined language understanding and generation capabilities, applied to more complex language understanding, text generation, fine-grained sentiment analysis, and other tasks; other models such as ELECTRA [8] and DeBERTa [9] have also shown

outstanding performance, achieving unprecedented results in multiple NLP benchmarks. These models optimize the training process and parameter efficiency, achieving better performance with fewer parameters, demonstrating the ongoing innovation and potential of Transformer models in the NLP field. Additionally, cross-language models, like mT5 [10] and XLM-R [11], utilizing Transformer technology, have greatly advanced research in machine translation and cross-cultural communication. All the aforementioned examples demonstrate the enormous potential of Transformer models in handling complex, large-scale datasets, and their key role in the future development of NLP.

Similar to NLP, in the field of CV, Transformer models have demonstrated their unique ability, particularly in processing spatial and temporal continuity in image data. This technology has been widely applied to various tasks such as image recognition [12], object detection [13], image segmentation [14], image super-resolution [15], video understanding [16], image generation [17], text-image synthesis [18], and visual question answering [19]. For example, models using Transformer technology can create unprecedentedly realistic and innovative images in image generation tasks [20], and in video understanding, they can identify and interpret dynamic images more accurately [21].

However, in the field of tactile technology, the application of Transformer models is still in its early stages. Tactile technology plays a crucial role in understanding the state and characteristics of objects, such as their shape, posture, texture, and material [22, 23, 24, 25, 26, 27]. Optical tactile sensors like GelSight [28], GelTip [29], TacTip [30], and DIGIT [31] offer higher-resolution tactile images compared to traditional single-point touch sensors and tactile arrays. These sensors are extensively used in various robot tasks, including sliding detection [32], texture recognition [33], surface tracking [34], object pushing [35], cable manipulation [36], insertion and tightening [37], and more.

Through these sensors, previous studies, such as [34, 35], often used tactile textures to extract handcrafted features, but these features were less effective when dealing with diversified and complex textured objects.

Recent research has taken a different approach by exploring end-to-end methods for learning tactile-motor strategies [32, 36, 37]. These methods directly extract features from tactile images using convolutional neural networks (CNNs), recurrent neural networks (RNNs), or Transformers, which then guide robot manipulation strategies. By enhancing the gradient consistency of the neural network, these methods improve the ability of tactile-motor strategies to generalize across different objects. Among them, approaches based on Transformer architectures have demonstrated significant promise. They excel in achieving robust and efficient tactile perception and manipulation within the field of robotics [38, 39].

In light of the advantages of Transformers in the tactile field, this review aims to comprehensively introduce the research progress in the tactile field adopting Transformer architecture, especially in applications with two types of inputs: Visuo-Tactile (a.k.a., vision-touch) and Only-Tactile (i.e., touch only). We will delve into various tasks in

the tactile field, such as object recognition, cross-modal generation, object manipulation, etc. (Sec. 4), and the effects of using Transformer-based models. Through this article, complemented by the structured overview provided in the mind map (Fig. 1), we aim to offer a detailed reference resource for researchers in the tactile field to understand and utilize Transformer architecture and to provide guidance for future research directions. Additionally, the article will discuss the challenges faced by Transformer architecture in the tactile field and the direction of future developments, as outlined in the thematic areas of our mind map.

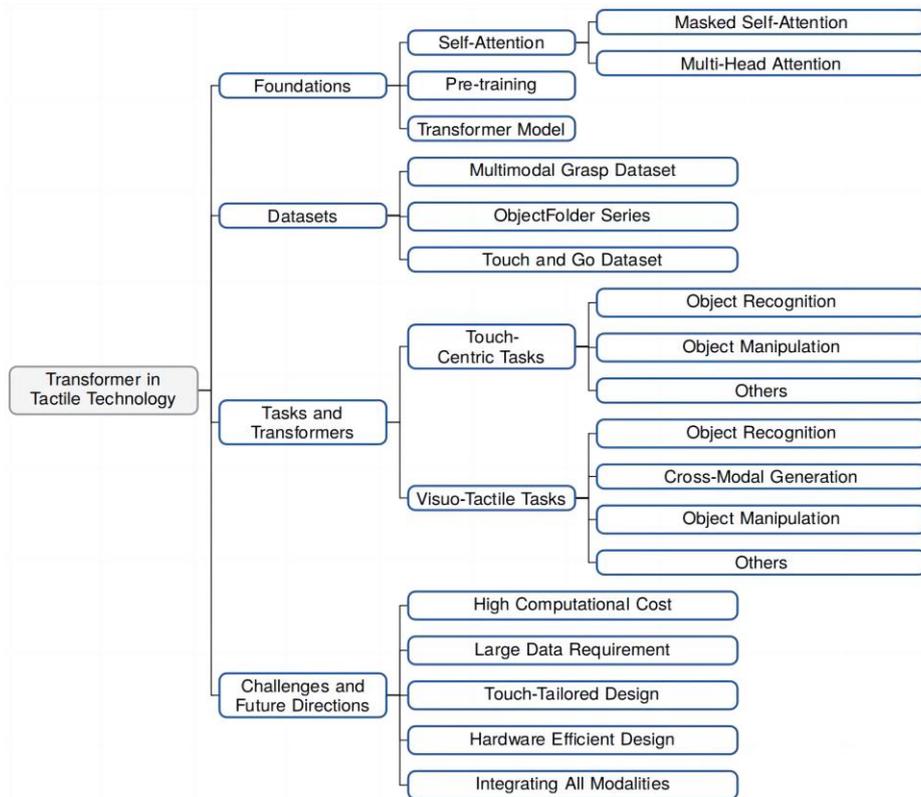


Fig. 1. Typology of Transformer in Tactile Technology

2 Foundations

In the context of developing Transformer models, two fundamental ideas significantly contribute to their evolution and progress [40]. First, the self-attention mechanism allows for capturing long-term dependencies between sequence elements, which is difficult for traditional recurrent models to achieve. Compared to recurrent models, the self-attention mechanism can encode these relationships more effectively. Second, pretraining is another core idea, referring to training on a large-scale labeled or unlabeled corpus in a self-supervised or supervised manner, followed by fine-tuning on relatively

small labeled datasets for target tasks [5, 41, 42]. These two ideas have played a vital role in the development of the Transformer model and have been widely applied in various Transformer networks.

Next, we will provide a brief introduction to these two key ideas (Sec. 2.1 and Sec. 2.2, respectively), and summarize Transformer networks that apply these ideas (Sec. 2.3). This prior knowledge will assist us in gaining a better understanding of the Transformer-based models utilized in the tactile field (Sec. 4).

2.1 Self-Attention in Transformers

For an item in a sequence, self-attention estimates the relevance of the item with other items (for example, which words in a sentence are likely to occur together). Self-attention is a core component of the Transformer, which is pivotal for modeling the interactions between all elements in a sequence. Essentially, a self-attention layer updates each item of the sequence by aggregating global information from the entire input sequence. Suppose a sequence contains n entities, represented as (x_1, x_2, \dots, x_n) , where d is the embedding dimension used to represent each entity. The goal of self-attention is to capture the interactions between all n entities by encoding global contextual information for each entity. This is achieved by defining three learnable weight matrices for transforming Queries ($W^Q \in \mathbb{R}^{d \times d_q}$), Keys ($W^K \in \mathbb{R}^{d \times d_k}$) and Values ($W^V \in \mathbb{R}^{d \times d_v}$). The input sequence X is first projected onto these weight matrices, obtaining $Q = XW^Q$, $K = XW^K$, and $V = XW^V$. The output $Z \in \mathbb{R}^{n \times d_v}$ of the self-attention layer is

$$\mathbf{Z} = \text{softmax}\left(\frac{\mathbf{QK}^T}{\sqrt{d_q}}\right)\mathbf{V}. \quad (1)$$

For a given entity in the sequence, self-attention first calculates the dot product of the query with all keys, followed by normalization with the softmax operator to obtain attention scores. Then, each entity becomes the weighted sum of all entities in the sequence, where weights are given by the attention scores.

Masked Self-Attention. The standard self-attention layer focuses on all entities. For the Transformer model [1], which is trained to predict the next entity in the sequence, the self-attention block used in the decoder is masked to prevent attention to subsequent future entities. This is simply accomplished by an element-wise multiplication operation with a mask $\mathbf{M} \in \mathbb{R}^{n \times n}$, where M is an upper triangular matrix. Masked self-attention is defined as

$$\text{softmax}\left(\frac{\mathbf{QK}^T}{\sqrt{d_q}} \circ \mathbf{M}\right) \quad (2)$$

where \circ represents the Hadamard product. Essentially, when predicting an entity in the sequence, masked self-attention sets the attention scores for future entities to zero.

Multi-Head Attention. To encapsulate the multiple complex relationships between different items in a sequence, multi-head attention contains multiple self-attention blocks (in the original Transformer model [1], the number of self-attention blocks $h = 8$). Each block has its own set of learnable weight matrices $\{\mathbf{W}^{Q_i}, \mathbf{W}^{K_i}, \mathbf{W}^{V_i}\}$, where $i = 0 \dots (h - 1)$. For the input X , the outputs of the h self-attention blocks in multi-head attention are then concatenated into a single matrix $[\mathbf{Z}_0, \mathbf{Z}_1, \dots, \mathbf{Z}_{h-1}] \in \mathbb{R}^{n \times h \cdot d_v}$ and projected onto a weight matrix $\mathbf{W} \in \mathbb{R}^{h \cdot d_v \times d}$.

The main difference between self-attention and convolution operations is that the filters are dynamically computed, rather than static filters in convolutions (which remain unchanged for any input). Additionally, self-attention is invariant to permutations and changes in the number of input points. Therefore, it can easily handle irregular inputs, while standard convolutions require a grid structure.

2.2 (Self) Supervised Pre-training

Transformer models based on self-attention typically employ a two-stage training mechanism. First, they undergo pre-training on large-scale datasets, which can be supervised or self-supervised. Subsequently, the pre-trained weights are adapted to small or medium-sized downstream task datasets.

Examples of downstream tasks in natural language processing include question answering, while in computer vision, tasks like image classification and object detection are common. In the fields of NLP and CV, the effectiveness of large-scale Transformer pre-training has been demonstrated. For instance, the Vision Transformer model (ViT-L) [12] saw a 13% absolute drop in accuracy on the ImageNet test set when trained only on the ImageNet training set (which has 50,000 images), compared to pre-training on the JFT dataset (which has 300 million images).

Since obtaining manual labels on a large scale is cumbersome, self-supervised learning has been very effectively used in the pre-training phase. Self-supervised pre-training plays a key role in enhancing the scalability and generalization abilities of Transformer networks, enabling them to train networks with over a trillion parameters (e.g., Google's Switch Transformer [43]).

Self-supervised learning offers a promising learning paradigm as it enables learning from large amounts of readily available unlabeled data. In the self-supervised pre-training phase, the model is trained to learn meaningful representations of underlying data by solving a pre-task [5]. The pseudo-labels of the pre-task are automatically generated based on data attributes and task definition (without any expensive manual annotation). Therefore, defining a pre-task in self-supervision is a key choice. Existing self-supervised methods can be broadly classified based on their pre-tasks into (a) generative methods, (b) context-based methods, and (c) cross-modal methods.

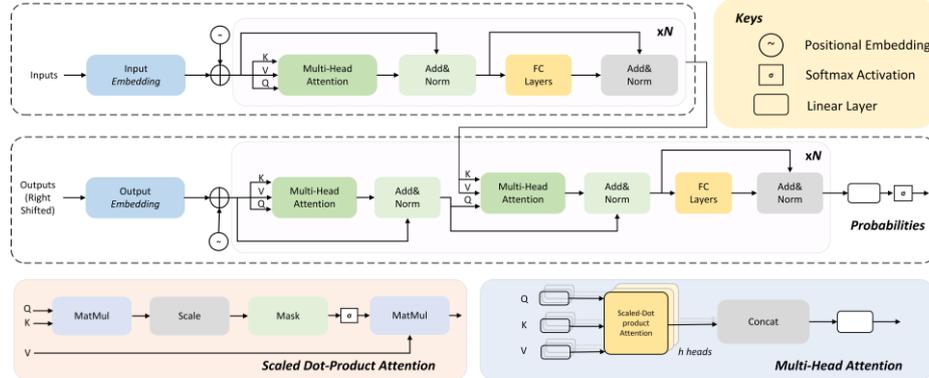


Fig. 2. The architecture of the Transformer model [1]. This model was originally developed for language translation tasks, where an input sequence in one language needs to be transformed into an output sequence in another language. The Transformer encoder (middle row) processes the input language sequence and converts it into embeddings before passing it on to the encoder modules. The Transformer decoder (bottom row) processes the previously generated outputs in the translation language and the encoded input sequences from the middle branch, to output the next word in the output sequence. The previously outputted sequence (used as the input to the decoder) is obtained by shifting the output sentence to the right by one position and adding a sentence-start token at the beginning. This shift prevents the model from simply learning to copy the decoder input to the output. The real data used to train the model is the output language sequence (without the right shift), with a sentence-end token appended. Blocks composed of multi-head attention (top row) and feed-forward layers are repeated N times in both the encoder and the decoder.

2.3 Transformer Model

The architecture of the Transformer model proposed in [1] is shown in Fig. 2. It has an encoder-decoder structure. The encoder (first row) consists of six identical modules (denoted as $N = 6$ in Fig. 2), each containing two sub-layers: a multi-head self-attention network, and a simple position-wise fully connected feed-forward network. After each module, residual connections [44] and layer normalization [45] are used.

It is noteworthy that, unlike conventional convolutional networks which perform feature aggregation and transformation simultaneously in the same step (e.g., through a convolutional layer followed by a non-linear operation), in the Transformer model, these two steps are decoupled, i.e., the self-attention layer only performs aggregation, while the feed-forward layer performs transformation. Similar to the encoder, the decoder in the Transformer model (bottom row) also contains six identical modules. Each decoder module has three sub-layers, the first two (multi-head self-attention and feed-forward) are similar to those in the encoder, while the third sub-layer performs multi-head attention on the output of the encoder modules.

The original Transformer model [1] was trained for machine translation tasks. The input to the encoder is a sequence of words (i.e., a sentence) in one language. Positional encoding is added to the input sequence to capture the relative positions of each word in

the sequence. The dimension of the positional encoding is the same as the input, $d = 512$, and can be learned or predefined, e.g., through sine or cosine functions. As an autoregressive model, the Transformer model's [1] decoder uses previous predictions to output the next word in the sequence. Therefore, the decoder takes input from both the encoder and the previous outputs to predict the next word in the translated language sentence. To facilitate residual connections, the output dimension of all layers remains the same, i.e., $d = 512$. The dimensions of the query, key, and value weight matrices in multi-head attention are set to $d_q = 64, d_k = 64, d_v = 64$.

3 Datasets

Tactile datasets generally originate from various sensing technologies and can be categorized based on their characteristics. These datasets can typically be characterized by sensing modes, sensing resolution, tactile representation, and application scenarios. Sensing modes refer to the types of tactile information captured by the dataset, such as force, temperature, or texture. Sensing resolution describes the dataset's ability to distinguish tactile details. Tactile representation involves how the dataset expresses tactile information, for example, through pressure maps, deformation maps, or other sensor data. Application scenarios are specific tactile applications for which the dataset can be used, such as robotic haptics, virtual reality, etc. Here are some common tactile datasets:

Multimodal Grasp Dataset [46]. This dataset is designed for robotic manipulation research, featuring a combination of visual and tactile data. It encompasses 2,550 sets of valid data, derived from nearly 3,000 test trials using 10 different objects. The dataset provides a detailed account of both successful and failed grasping attempts, with a total success rate of 66.27%. It includes tactile data from 16-channel sensors across 400 steps in 24 seconds, along with visual data consisting of four images per grasp from two different camera viewpoints. Additionally, it captures the position sequence of eight motor joints and contains video data comprising two RGB and one depth video for each grasp. This rich dataset serves as a comprehensive tool for understanding the nuances of robotic grasping and manipulation, offering valuable insights into slip detection, and object classification.

ObjectFolder series datasets

- ObjectFolder 1.0 [47]: Introduced in 2021, this dataset includes 100 virtualized objects, each encoded with visual, auditory, and tactile sensory data. It's primarily focused on multisensory object recognition tasks and uses a uniform, object-centric, and implicit representation for each object's visual textures, acoustic simulations, and tactile readings. This makes the dataset flexible for use in various applications like instance recognition, cross-sensory retrieval, 3D reconstruction, and robotic grasping. However, it is noted for its limited scale and the quality of multisensory data, which might affect its generalization to real-world scenarios.

Table 1. Comparative Summary of Different Datasets Used in Tactile Domain Tasks. Each dataset is evaluated based on several parameters such as the number of object instances, total touches, data source, application in real-world scenarios, the environment in which data was collected, and the types of sensors used.

Dataset	Object inst.	Touches	Source	Real-world	Environment	Sensor
Multimodal Grasp [46]	3971	2550	Robot	✓	Tabletop	Eagle Shoal Hand
ObjectFolder 1 [47]	100	N/A	Synthetic	✗	N/A	DIGIT [31]
ObjectFolder 2 [48]	1000	N/A	Synthetic	✗	N/A	GelSight [28]
ObjectFolder-Real [49]	100	N/A	Synthetic	✓	Indoor	GelSight [28]
Touch and Go [50]	3971	13.9k	Human	✓	Indoor	GelSight [28]

- **ObjectFolder 2.0 [48]:** Launched in 2022, this dataset expands significantly on its predecessor. It features 1,000 multisensory neural objects, offering a tenfold increase in the number of objects and substantial improvements in rendering time and multisensory data quality. These enhancements aim to address the limitations of OBJECTFOLDER 1.0, particularly regarding scale and data quality. ObjectFolder 2.0 is used to facilitate research in multisensory learning, particularly in computer vision and robotics, and supports tasks like object scale estimation, contact localization, and shape reconstruction.
- **ObjectFolder-Real [49]:** This dataset contains multisensory measurements for 100 real-world household objects. It builds upon a newly designed pipeline for collecting 3D meshes, videos, impact sounds, and tactile readings of these objects. The focus here is on providing a realistic and practical representation of objects, incorporating data collection processes tailored for each sensory modality. This dataset plays a crucial role in multisensory object-centric learning, particularly in tasks like object recognition, reconstruction, and manipulation.

Touch and Go Dataset [50]. This dataset is a rich collection of natural vision-and-touch signals, featuring approximately 13.9k detected touches and around 3971 individual object instances. Gathered by human collectors in diverse environments, it encompasses a broad spectrum of objects, both rigid and deformable, located in diverse settings that include indoor spaces like university buildings and apartments, as well as outdoor environments such as hiking trails and playgrounds. A significant feature of this dataset is its use of the GelSight tactile sensor, which records detailed tactile and visual information. The dataset's diversity, covering a broad range of real-world scenes and object types, makes it particularly valuable for research in multimodal learning. It's further enhanced by annotations for material categories and specific frame tagging, which are crucial for in-depth studies of visuo-tactile interactions in natural environments. *错误! 未找到引用源。*

The section effectively summarizes various tactile datasets, each tailored for specific research applications in robotic manipulation and multimodal learning. These datasets, detailed in Table 1, range from the Multimodal Grasp Dataset with its focus on robotic manipulation, to the ObjectFolder series that progresses from virtualized objects to real-world scenarios, and the Touch and Go Dataset that captures natural vision-and-touch signals in diverse environments. Each dataset contributes unique insights and data, enriching the field of tactile sensing and offering valuable resources for research and development in this area.

4 Tasks and Transformers

4.1 Touch-Centric Tasks and Tactile-Only Transformers

Firstly, we review tasks and the corresponding Transformer-based methods that only use tactile data.

4.1.1 Object Recognition

For the task of object surface recognition, Awan et al. [51] proposed a novel tactile texture classification method based on the Transformer architecture. In their study, the authors collected acceleration signals from 9 samples, which, after preprocessing and normalization, were segmented into 500 sample windows in the time domain as inputs to the model. After positional encoding, these were fed into three encoder layers, each with three attention heads. Notably, the model built by the authors only deployed the encoder module of the transformer, without using the decoder module. Each encoder layer is identical, containing two sub-layers: a self-attention sub-layer and a feedforward sub-layer. The feedforward sub-layer uses 1D-CNN and ReLU as the activation function, followed by a fully connected layer. Additionally, a normalization layer follows each sub-layer. For predictions, the output of the encoder module is fed into a Dense layer, followed by a softmax layer. Ultimately, the model proposed by the authors achieved an average accuracy of 98.87% and an F1 score of 96.89%, reaching state-of-the-art classification accuracy in texture classification tasks.

4.1.2 Object Manipulation

For the object manipulation task in grasping, non-destructive grasping of deformable objects has always been a challenge in the field of robotics. Humans can perceive the physical properties of objects and apply precise forces while grasping fragile targets to accomplish dexterous and non-destructive operations. To transfer this capability from humans to robots, an improved Transformer model for processing complete tactile time series during the grasping process was proposed in [52]. Also, considering the scarcity of grasping tactile datasets available in the tactile domain compared to computer vision, the authors also established a tactile dataset of 9375 grasps of 15 types of fruits. The proposed model achieved a fruit recognition accuracy of 97.33%

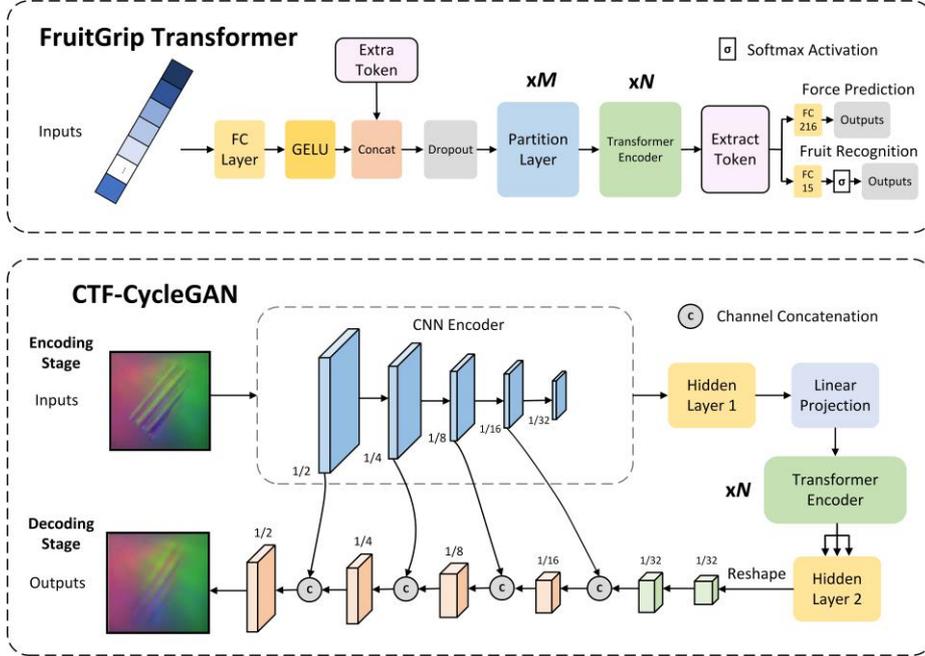


Fig. 3. Structural Display of the FruitGrip Transformer and CTF-CycleGAN Models for Object Manipulation Tasks.

this dataset, surpassing traditional recurrent neural network (RNN) models. The results suggest that this work could facilitate tactile perception and non-destructive grasping in robots in agricultural fields, aiding in fruit picking, processing, sorting, and other related areas. This method also provides a technical reference for other studies on grasping tactile data.

For the task of manipulating tubular objects like test tubes, developing robots capable of handling them has the potential to accelerate experimental processes. However, for robots, learning to process tubular objects with a single perception and bridging the gap between simulation and reality remains a challenge. For this purpose, Zhao et al. [53] proposed a Sim2Real transferable hand posture estimation method based on tactile images, specifically for manipulating tubular objects. The core of this method is the CTF-CycleGAN model, a novel pixel-level unsupervised domain adaptation network that cascades CNNs and Transformers, aimed at narrowing the pixel-level domain differences in tactile tasks by introducing attention mechanisms and task-related constraints. Eventually, this method was implemented in a reinforcement learning-based policy learning framework for learning Sim2Real transferable robot insertion and removal actions, further applied in human-robot collaborative test tube placement scenarios and robotic pipetting scenes. Experimental results showed that the learned tactile-motor strategy has generalization capabilities for tubular object manipulation.

To better illustrate the FruitGrip Transformer and CTF-CycleGAN models, Fig. 3 depicts their structural designs. As shown, the FruitGrip Transformer model uses a framework similar to the Visual Transformer (ViT) but with specific adaptations for tactile data. The model directly projects one-dimensional tactile sequences of 216-dimensional vectors to a specified dimension and then processes them using a Transformer encoder. Uniquely, it segments the input sequence regionally, simulating tactile perception in different parts of a human hand, and then processes it through Transformer blocks of various scales. This structure helps the model effectively learn and predict the type and grasping force of fruits, facilitating dexterous manipulation of deformable objects.

In contrast, the CTF-CycleGAN model combines Convolutional Neural Networks (CNNs) and Transformers in an encoder-decoder structure. In the encoding phase, it first uses multiple convolutional layers to extract high-level features from tactile images, then enhances feature extraction capabilities with a self-attention Transformer encoder. Its distinguishing feature is the ability to capture global information in tactile images, focusing on regions with significant features. The decoding phase uses a U-shaped structure to combine features from the Transformer encoder and CNNs for image generation.

In summary, the FruitGrip Transformer is primarily focused on precise object recognition and force prediction through the analysis of tactile time series data, while the CTF-CycleGAN is dedicated to the generation of high-quality images and advanced feature extraction, with a specific emphasis on tactile imagery processing. Although both models incorporate Transformer technology, they are distinct in their application objectives and structural designs, showcasing the versatility and adaptability of deep learning in diverse robotic applications.

4.1.3 Others

Compared to other multimedia applications like videos, sounds, and images, the transmission of tactile signals requires higher reliability and lower latency. Vibrational tactile signals can represent key elements of rich tactile information, such as textured surfaces. To transmit these signals to remote locations, and create immersive and realistic user experiences, the development of an efficient tactile compression method is needed. In [54], a Transformer-based vibrational tactile signal codec TVSC was proposed, combining Transformer structure, quantization, and entropy coding. It is trained end-to-end for the encoder and decoder to achieve efficient compression of tactile signals. After inverse quantization, the decoder is fine-tuned to overcome the impact of signal loss caused by quantization. When the compression ratio is small, the signals reconstructed by TVSC closely match the original signals. As the compression ratio increases, the quality of the reconstructed signals slightly decreases and stabilizes. Experimental results show that TVSC outperforms PVC-SLP [55].

4.2 Visuo-Tactile Tasks and Visuo-Tactile Transformers

Tasks that combine sight and touch involve the combination of visual and tactile sensory inputs to carry out different functions or accomplish specific goals. These tasks led to the development of Visuo-Tactile Transformers.

We review tasks and the corresponding Transformer-based methods that integrate visual and tactile sensory modes.

4.2.1 Object Recognition

In the field of computer vision, there has been extensive research and application of object recognition technology. Specifically, 2D and 3D images have been effectively used for classifying and recognizing objects and their surface properties. For instance, convolutional neural networks (CNNs) have been successfully applied to signal recognition processes [56]. Moreover, the application of one-dimensional CNNs and bidirectional long short-term memory networks (Bi-LSTMs) in-vehicle network anomaly detection was demonstrated in [57].

Unlike touch-based methods, vision-based methods require the use of optical sensors, which may be limited by obstructions and lighting conditions [58]. Conversely, methods that combine both visual and tactile modalities have been proven to outperform methods using only the visual modality [59].

However, integrating visual and tactile features faces significant challenges. Since visual and tactile features exist in different feature spaces, it can be challenging for fusion modules to establish direct interactions between them. Current methods handle visual and tactile features during the fusion stage but struggle to bridge the significant gap between these two modalities. For instance, visual data and tactile data vary in terms of frequency, signal type, and perceptual domain. More specifically, visual sensors like cameras collect data through low-frequency instantaneous captures, while tactile sensors acquire data through high-frequency continuous sliding or touching. The signal type of visual data consists of RGB images composed of pixels, while tactile data includes various signal types, such as vibration signals, force feedback, and tactile images. Regarding the perceptual domain, visual perception focuses on the perception of shape and local color, while tactile perception emphasizes pressure and roughness. All these differences make it challenging to directly fuse information from these two modalities.

To address the aforementioned issues, Wei et al. [60] proposed a new visual-tactile Transformer-based method based on Alignment and Multi-Scale Fusion (AMSF). First, a visual-tactile contrastive learning module is used to match visual and tactile data, and then the matched features are fed into a multi-scale fusion module through a Transformer, achieving deep interaction between the two modalities. By leveraging these two modules, this method not only aligns visual and tactile modality but also

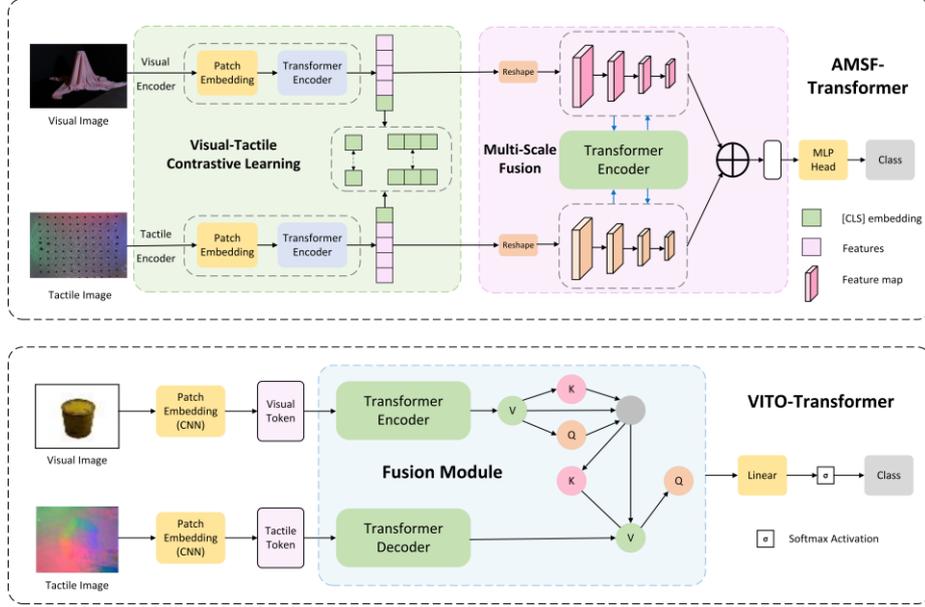


Fig. 4. Structural Display of the AMSF-Transformer and VITO-Transformer Models in Object Manipulation Tasks.

achieves a fusion of multi-scale feature maps. Extensive experiments on three public datasets (GelFabric [28], LMT [61, 62], and Vitac [22]) demonstrated the superiority of the proposed method. Subsequently, Li et al. [63], to address the challenge of fusing visually and tactilely information, designed a visual and tactile information fusion mechanism based on the Transformer architecture, named VITO-Transformer. Thanks to this fusion mechanism, the accuracy of object recognition was significantly improved. In the end, comprehensive comparative experiments were conducted on the public visual-tactile dataset (ObjectFolder 2.0) and a custom visual-tactile dataset (700 object visual images and 7 object tactile images), confirming the advantages of VITO-Transformer and validating the effectiveness of this fusion mechanism compared to current popular fusion algorithms, bringing innovative solutions to the field of visual-tactile fusion.

To delve deeper into and compare the AMSF-Transformer with the VITO-Transformer, Fig. 4 presents their structural details. AMSF utilizes a contrastive learning approach to align visual and tactile information, followed by multi-scale fusion using a Transformer. The crux of this model lies in using a visual Transformer as a dual-modal encoder to process input visual and tactile images. The alignment of these two representations is achieved through visual-tactile contrastive loss. Its multi-scale fusion module operates at every layer of the multi-modal encoder, ensuring effective feature fusion, and ultimately processes the fused features through an MLP for classification.

In contrast, The VITO-Transformer model adopts an innovative decoding mechanism for the fusion of visual and tactile data. The model initially segments and transforms

the input visual image into token embedding sequences, which are then fed into the Transformer encoder. Additionally, tactile data is converted into tactile image data and similarly segmented and transformed into token embedding sequences. These sequences, processed by the Transformer encoder, are subsequently input into the decoder. The VITO-Transformer's decoder crucially integrates visual and tactile information by adding an extra layer of multi-head attention operations. This unique mechanism allows the decoder to combine and process these two different types of information effectively, generating fused feature information for precise identification of various objects.

Although both models are dedicated to the fusion of visual and tactile information, they differ in their implementation methods and structures. AMSF focuses on information alignment and layered fusion, while VITO-Transformer emphasizes information fusion through its specialized decoder mechanism. This comparison reveals the diverse approaches to multi-modal perception in the field of machine learning, providing valuable insights for future research.

In the field of autonomous driving, to achieve accurate road recognition (object recognition task for identifying specific roads), a multimodal fusion recognition network based on the CNN-Transformer structure was proposed in [64]. Specifically, visual and tactile channels were input into specific channel SE-CNNs, which emphasized valuable input information to obtain weighted features. These features were then input into a "bottleneck" fusion Transformer encoder and produced recognition results. Additionally, the authors designed a fusion feature extractor to enhance fusion representation capabilities and improve fusion accuracy. To test the model's performance, the authors conducted vehicular field experiments and established a dataset composed of four types of road surfaces, demonstrating that the network achieved an accuracy of 99.48% in road recognition tasks, showing its strong application value.

4.2.2 Cross-Modal Generation

For cross-modal generation tasks, fully capturing the underlying correlations between multimodal data and generating more accurate tactile data for high-fidelity tactile rendering remains a challenging issue. To address this, a vision-audio assisted tactile Transformer model for cross-modal generation of tactile friction coefficients was proposed in [65]. The model first encodes the time-frequency graph of audio and the RGB image input of vision to extract joint encoding features. Then, these joint encoding features are transformed into tactile decoding features by a Transformer module. After that, the tactile decoding features are used to reconstruct tactile amplitude and phase spectra. Finally, they are converted into 1-D friction coefficients through Inverse Short Time Fourier Transform (ISTFT) [66]. This method combines self-attention mechanisms and convolution operations to more comprehensively capture the correlations between audio-visual-tactile data. Ultimately, the model was evaluated on the LMT tactile materials dataset [56]. Results showed that the proposed tactile Transformer model outperformed AVHR [67] and convolutional autoencoder (CAE) methods [68], improving Root Mean Square Error (RMSE) by 28.53% and 33.25%, respectively. Comparative experimental results also indicated that lighting conditions affect the accuracy of the

vision-tactile model, with RMSE for low-light images being 24.10% higher than for normal images, but the combination of audio and low-light images could compensate for errors caused by lighting conditions. Additionally, the tactile Transformer model proposed by the authors achieved a tactile rendering fidelity of 92.3% compared to AVHR and CAE methods, significantly enhancing rendering fidelity ($p < 0.05$). This result was obtained through a specific statistical significance test, namely the Analysis of Variance (ANOVA). However, the ability of the model to generate materials for new categories could not be fully verified due to the unbalanced distribution of different material categories in the existing datasets.

For specific domain cross-modal generation, models combining visual and tactile modalities based on the Transformer architecture also play an important role. For example, in the VR/AR field, a method for generating tactile signals based on the Transformer network framework was proposed in [69]. This method uses the visual image of an object's surface as visual data and scanning parameters (scanning speed v , position p , and normal force f) generated by a pen sliding over the object's surface as acceleration signals. The acceleration signals are synthesized through a Transformer-based generation model and a multimodal feature embedding module as the tactile data for the object's surface. Unlike traditional Transformer structures applied in NLP, this model replaces the input/output embedding layers with linear layers to adapt to the size of the encoder/decoder inputs and uses a linear projection layer to reshape the final output of the decoder to the target size. Especially, this model also adds self-attention layers, multimodal embedding layers, and feedforward layers in the latent space to adapt to the multimodal fusion of visual-tactile data. To evaluate the model's performance, the authors used the aluminum texture as a test sample and employed a sliding window-based data augmentation strategy to create a multimodal dataset. Finally, the Root Mean Square Error (RMSE) value for the selected samples was 0.0082 (SD = 0.0023). In the future, the authors will conduct user studies to assess the realism of the generated virtual tactile textures and further integrate the tactile texture modeling and rendering framework with VR/AR devices to improve the user experience of texture simulation in virtual environments.

4.2.3 Object Manipulation

The object manipulation task is subdivided into four sub-tasks: 1) slip detection, 2) object grasping, 3) object rotation, and 4) combined manipulation of objects. Each sub-task's visuo-tactile Transformer architecture models will be introduced in detail.

In object manipulation tasks, slip detection plays a crucial role in dexterous robot gripping and manipulation and has always been a challenging problem in the field of robotics. However, due to limitations in perception capabilities, the desired detection performance may not be achievable through tactile perception alone. Adding visual perception is an intuitive solution. Li et al. [70] proposed a vision-tactile learning method based on deep neural networks to detect slip, highlighting the importance of vision-tactile fusion perception in slip detection tasks. Compact and multimodal representations of sensory inputs learned through self-supervised learning can be used to improve

the sample efficiency of policy learning [71]. However, these early fusion methods, which directly combine features of the visual and tactile modalities, have been shown to lack the ability to capture modality-specific and cross-modal features [72]. Furthermore, most of the methods mentioned above are explicitly designed for tactile sensors with different data formats and sampling rates, leading to significant limitations in platform generalizability.

To address the inconsistency of vision-tactile sequence learning in slip detection tasks, Cui et al. [73] proposed a generalized vision-tactile Transformer (GVT-Transformer) model. Specifically, the GVT-Transformer mainly consists of modality-specific and cross-modality Transformers, capturing features of specific modalities and cross-modality, respectively. Experimental results showed that this method not only outperformed traditional early fusion methods in detection performance but also was more suitable for non-aligned situations. Slip detection experiments conducted on real robots also confirmed the Sim2Real generalization capability of the proposed method.

In the object grasping task, safe grasping of deformable objects remains a challenging task. Loose grasps with small grasping forces can lead to object slip, while large grasping forces can cause object damage. Additionally, the contact geometry shape and friction characteristics of an object might affect the optimal grasping force for safe grasping. To address this issue, a generalizable vision-tactile robotic grasping method based on the Transformer architecture was proposed in [37]. In this method, the Transformer model performs two predefined exploration actions (pinching and sliding), not only predicts the outcome of various grasping intensities using a multilayer perceptron (MLP) but also enhances its grasp prediction capabilities by learning physical feature embeddings derived from sensor feedback. Using these predictions, the gripper can infer safe grasping intensities. To evaluate the model's performance, the authors first conducted benchmark tests for slip detection tasks on a public dataset containing 84 everyday objects. The results showed that the model far surpassed the CNN+LSTM model in terms of grasping accuracy and computational efficiency. Besides, the authors collected a new dataset containing six types of fruits (i.e., plums, oranges, lemons, tomatoes, apples, and kiwifruits) for grasping, and used the model for online grasping experiments on both seen and unseen fruits. Finally, the authors extended the model to different-shaped objects and demonstrated the effectiveness of the pre-trained model trained on a large-scale fruit dataset.

In the object rotation task, accurate control of robot finger movements is required, and tactile feedback must be received and processed in real time to adjust rotation strategies. For this purpose, a vision-tactile Transformer-based model RotateIt was proposed in [74]. This model effectively processes multimodal perceptual information by integrating visual information provided by depth cameras and tactile data from tactile sensors. The authors first built a dataset of objects with different physical properties and random initial postures from EGAD, Google Scanned Objects, YCB, and ContactDB, with a width/depth/height ratio of less than 2.0. Visual information is processed by a convolutional neural network (ConvNet), while tactile information is transformed into a discrete contact point position representation. The RotateIt model combines these visual

and tactile features to achieve multi-axial object rotation based on fingertip. Experimental results showed that RotateIt could effectively understand the geometric shape and contact point characteristics of objects, enabling robots to manipulate objects more naturally and effectively.

In the task of combined manipulation of objects, existing vision-tactile methods typically use latent representations that are vectors or clusters in R^n and do not utilize the attention architecture [1]. These representations might flatten details and face difficulties in the locality of the images. These problems are particularly tricky for manipulation, as the difference between contact and non-contact might be small in pixel space, or the robot might need significant movement to make contact with an object. To address this issue, Chen et al. [75] proposed a model for object manipulation, the Visuo-Tactile Transformer, abbreviated as VTT. This model uses self-attention mechanisms and cross-modal attention mechanisms to fuse visual and tactile information and improves feature learning through learned embedding vectors. The authors also proposed three types of learnable embeddings to enhance model performance. Finally, the model was combined with reinforcement learning to form a stable and robust manipulation method. The model was experimentally evaluated on four simulated robotic tasks (Pushing, Door-Open, Picking, Peg-Insertion) and a real-world pushing task. Compared to baseline methods such as concatenation and product of expert (PoE), VTT demonstrated effectiveness in representation learning.

4.2.4 Others

In terms of multimodal representation, Noguchi et al. [76] proposed a model that learns visual, tactile, and proprioceptive perception to obtain multimodal plastic body and near-personal space representations. This approach includes three main stages: 1) encoding of visual, tactile, and proprioceptive perceptions, 2) integration of encoded perceptions through Transformer architecture utilizing a self-attention mechanism, 3) reconstruction of visual and tactile perceptions from integrated sensory representations. The model, focused on a simulated robot arm, facilitates the development of multimodal and body-centered peripersonal space representation through tool use, surpassing previous models that offered non-body-centered plastic body representations. By learning from camera vision, arm touch, and proprioception of camera and arm postures, it achieves a body representation that localizes tactile sensations on a concurrently developed peripersonal space representation. This learning process, especially during tool use, endows the body representation with

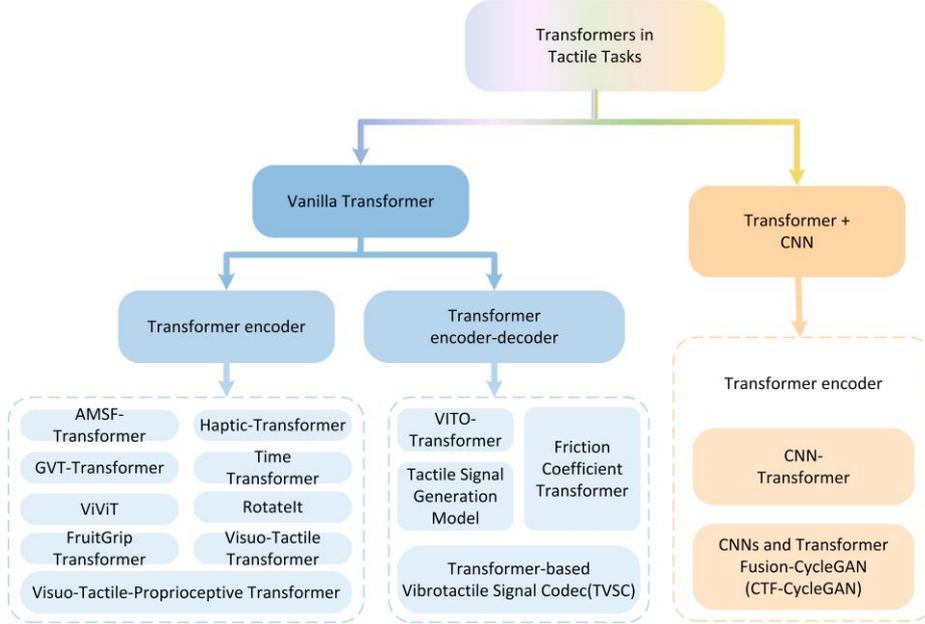


Fig. 5. Taxonomy of Transformer-based Approaches for Tactile Tasks: First divided into Vanilla Transformer and Transformer + CNN, then further subdivided into Transformer encoder and Transformer encoder-decoder.

plasticity, allowing it to adapt to changes in interaction brought by tool use while maintaining an internal spatial coordinate system. The research advances the understanding of multimodal representation development and the plasticity of self-body cognition, providing insights into cognitive neuroscience and developmental robotics.

5 Open Challenges and Future Directions

Despite the outstanding performance and intriguing features of Transformer models (as discussed in Table 2 and Fig. 5), applying them in practical settings still presents several challenges. The most significant bottlenecks include the need for large amounts of training data and the associated high computational costs. There are also challenges in visualizing and interpreting Transformer models. In this section, we outline these challenges, mention some recent efforts made to address these limitations and highlight open research questions.

5.1 High Computational Cost

As discussed in Sec.1, a key advantage of Transformer models is their flexibility to scale up to high parameter complexity. While this is a notable feature, allowing for

Table 2. Comprehensive Overview of Transformer-based Approaches for Tactile Tasks, Including Method Names, Associated Tasks, Utilized Datasets, Evaluation Metrics, and Key Design Highlights.

Method	Task	Datasets	Metrics	Design Highlights
AMSF-Transformer [38]	Object Recognition	GelFabric, LMT, Vitac	Accuracy, Precision	New alignment strategy, multimodal contrastive learning.
VITO-Transformer [6338]	Object Recognition	OBJECTFOLDER 2.0, Self-made Dataset	Accuracy, Precision, Recall and F1	Visual-tactile fusion via Transformer network.
CNN-Transformer [6438]	Object Recognition	Four different kinds of roads Dataset	Accuracy	Multimodal fusion with CNN-Transformer, channel-specific feature weighting.
Haptic-Transformer [5138]	Object Recognition	Fabric, steel, plastic and wood Dataset	Accuracy, Precision, Recall and F1	Transformer model for tactile texture classification.
Friction Coefficient Transformer [6538]	Cross-Modal Generation	LMT Haptic Material Dataset	RMSE, Relative Error	Self-attention for audio-visual encoding, tactile friction generation.
Tactile Signal Generation Model [6938]	Cross-Modal Generation	Enhanced HaTT Dataset	RMSE	Visual-tactile input to simulate tactile textures in AR/VR.
Generalized Visual-Tactile (GVT) Transformer [7338]	Object Manipulation	Visual-Tactile Dataset	Precision, Recall and F1	Modality-specific and cross-modality Transformers for non-aligned situations.
TimeTransformer and ViViT [3738]	Object Manipulation	Slip Detection: 84 daily objects. Safe Fruit Grasping: Involving six different types of fruits.	Accuracy, Success Rate	Self-attention for deformable object tasks, outperforms CNN+LSTM.
RotateIt [7438]	Object Manipulation	EGAD, Google Objects, YCB, ContactDB	TTF, RotR, RotP and Rotations	Vision-tactile integration, depth cameras and tactile sensors fusion.
Visuo-Tactile Transformer (VTT) [7538]	Object Manipulation	RGB images, wrist sensor data	Success Rate	Self and cross-modal attention, learnable embeddings for feature enhancement.
FruitGrip Transformer [5238]	Object Manipulation	Fruit grasping dataset	Accuracy, Speed and Parameters	Processes tactile time series, tactile dataset establishment.
CTF-CycleGAN [5338]	Object Manipulation	Real-world tactile images Dataset	Success Rate	Pixel-level domain adaptation, CNN-Transformer cascade, attention mechanisms.
Visuo-Tactile-Proprioceptive Transformer [7638]	Others	Visuo-tactile-proprioceptive Dataset	PC space visualization	Body representation model with tool-use-induced plasticity.

Transformer-based Vibrotactile Signal Codec (TVSC) [5438]	Others	TUM Dataset and LMT tactile material Dataset	PSNR, Compression Ratio	Vibrotactile signal compression using Transformer-based codec.
---	--------	--	-------------------------	--

the training of enormous models, it leads to higher training and inference costs. For instance, the training time for the base BERT model [5] (with 110 million parameters) is approximately 1.89 petaflop days, while the GPT-3 model [6] (with 175 billion parameters) requires an astonishing ~ 1925 times increase, approximately 3640 petaflop days. This comes with a hefty cost tag; for example, according to an estimate [77], GPT-3's training could have cost OpenAI 4.6 million. Furthermore, these large-scale models require significant compression (such as distillation) to be feasible in real-world settings.

5.2 Large Data Requirement

Because Transformer architectures do not inherently encode inductive biases (i.e., prior knowledge) for processing tactile data, they typically require extensive training to figure out the modality-specific underlying rules. For instance, CNNs have built-in translation invariance, weight sharing, and local scale invariance due to pooling operations or multi-scale processing blocks. However, Transformer networks need to figure out these tactile data-specific concepts from training examples on their own, leading to longer training times, a significant increase in computational demands, and the need for large datasets to process.

5.3 Touch-Tailored Transformer Design

It is observed that the majority of previous studies in the field of tactile tasks tend to directly utilize NLP Transformer models for addressing problems in the tactile domain. These include architectures designed for object recognition, object manipulation and especially multi-modal processing. While the preliminary results of these straightforward applications are quite encouraging and motivate further exploration of the advantages of self-attention and self-supervised learning, the current architectures may still be more suited to language problems (using sequential structures) and require further intuition to make them more effectively applicable to tactile inputs.

5.4 Hardware Efficient Design

Large Transformer networks can have dense power and computational demands, hindering their deployment on edge devices and in resource-constrained environments (such as internet-of-things platforms).

This challenge becomes even more critical in the tactile field, where the need for real-time haptic feedback imposes stringent requirements on latency and processing speed. The Tactile Transformer, which facilitates the transmission of touch and manipulation

in real-time over digital networks, demands not only low latency but also high reliability and rapid processing to simulate real-world physical interactions accurately.

Recent reports suggest efforts in compressing and accelerating NLP models on embedded systems like Field-Programmable Gate Arrays (FPGAs) [78]. Li et al. [78] used an enhanced block circulant matrix-based representation to compress NLP models and proposed a new FPGA architecture design to efficiently manage resources for high throughput and low latency. They were able to achieve 27-, 3-, and 81-times performance improvement (measured in FPS throughput) over the RoBERTa model [41] on CPU, reducing power consumption and energy efficiency.

To achieve this goal, Wang et al. [79] proposed using a neural architecture search strategy to design Hardware Aware Transformers (HAT) [80, 81, 82]. Specifically, a Super-Transformer model [79] is first trained as a performance surrogate, capable of estimating model performance without full training. This model contains the largest possible model in the search space and shares weights across common parts. Ultimately, an evolutionary search is performed, considering hardware latency constraints, to find suitable SubTransformer models [79] for the target hardware platform (e.g., internet-of-things (IoT) devices, GPU, CPU). However, such hardware-efficient designs are currently lacking, to enable seamless deployment of Tactile Transformers on resource-limited devices. Moreover, the search cost of evolutionary algorithms remains substantial due to the associated environmental impact of CO₂ emissions.

5.5 Towards Integrating All Modalities

Given that transformers provide a unified design to process different modalities, recent efforts have also focused on proposing more general-purpose, transformer-based universal reasoning systems. Inspired by biological systems capable of processing different forms of information, the Perceiver model [83] aims to learn a unified model that can handle any given input modality without making specific domain structural assumptions. The perceiver employs an asymmetric cross-attention approach to distill input information into low-dimensional latent bottleneck features to scale to high-dimensional inputs. Once features are refined into a compact and fixed-dimensional form, regular Transformer blocks are applied in the latent space. The initial Perceiver model demonstrated performance competitive with ResNets [84] and ViTs [12] in image classification and could handle 3D data, audio, images, videos, or their combinations. However, the model could only generate fixed outputs like class probabilities.

A recent improvement, termed Perceiver IO [85], aims to learn models with flexible inputs and arbitrary-sized outputs. This allows applications to problems requiring structured outputs, like natural language tasks and visual understanding. While these models avoid reliance on modality-specific architectural choices, the learning itself still involves modality-dependent choices, such as specific augmentations or positional encodings. An interesting and open direction for the future is to achieve complete modality agnosticism.

6 Conclusion

This review begins with a detailed introduction to the two fundamental ideas of the Transformer model and its traditional architecture. It then delves into the application of the Transformer architecture in various tactile tasks, such as object recognition, cross-modal generation, and object manipulation. The approaches discussed mainly fall into two categories: models that rely solely on tactile data and those that integrate tactile and visual data.

Our exploration encompassed not only the architectural innovations and their applications but also the challenges and potential future developments. Key challenges such as the high computational cost, large data requirements, and the need for touch-tailored Transformer designs were discussed, providing a comprehensive overview of the current state and prospects of Transformer models in the tactile domain.

We hope that this effort will further stimulate interest in the tactile community to harness the potential of Transformer models and to improve their current limitations.

Acknowledgments. This work is supported by the Talent Fund of Beijing Jiaotong University (2023XKRC006) and the Pattern Recognition Center, WeChat AI, Tencent Inc.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

7 References

1. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
2. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al.: "Huggingface's transformers: State-of-the-art natural language processing". *arXiv preprint arXiv:1910.03771* (2019)
3. Wei, X.; Cui, X.; Cheng, N.; Wang, X.; Zhang, X.; Huang, S.; Xie, P.; Xu, J.; Chen, Y.; Zhang, M.; Jiang, Y.; Han, W.: "Zero-Shot Information Extraction via Chatting with ChatGPT". *arXiv e-print arXiv:2302.10205* (2023)
4. Wang, P.; Wei, X.; Hu, F.; Han, W.: "TransGPT: Multi-modal Generative Pre-trained Transformer for Transportation". *arXiv e-print arXiv:2402.07233* (2024)
5. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K.: "Bert: Pre-training of deep bidirectional transformers for language understanding". *arXiv preprint arXiv:1810.04805* (2018)
6. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al.: Language models are few-shot learners. In: *Advances in Neural Information Processing Systems*, 33, pp. 1877–1901 (2020)
7. Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y.T.; Li, Y.; Lundberg, S.; et al.: Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023)
8. Clark, K.; Luong, M.-T.; Le, Q.V.; Manning, C.D.: Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555* (2020)

9. He, P.; Liu, X.; Gao, J.; Chen, W.: DeBERTa: Decoding-enhanced bert with disentangled attention. arXiv preprint arXiv:2006.03654 (2020)
10. Xue, L.; Constant, N.; Roberts, A.; Kale, M.; Al-Rfou, R.; Siddhant, A.; Barua, A.; Raffel, C.: mT5: A massively multilingual pre-trained text-to-text transformer. arXiv preprint arXiv:2010.11934 (2020)
11. Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; Stoyanov, V.: Unsupervised cross-lingual representation learning at scale. arXiv preprint arXiv:1911.02116 (2019)
12. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
13. Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision, pp. 213–229. Springer (2020)
14. Ye, L.; Rochan, M.; Liu, Z.; Wang, Y.: Cross-modal self-attention network for referring image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 10502–10511 (2019)
15. Yang, F.; Yang, H.; Fu, J.; Lu, H.; Guo, B.: Learning texture transformer network for image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 5791–5800 (2020)
16. Sun, C.; Myers, A.; Vondrick, C.; Murphy, K.; Schmid, C.: "Videobert: A joint model for video and language representation learning". In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 7464–7473 (2019)
17. Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; Gao, W.: Pre-trained image processing transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 12299–12310 (2021)
18. Reddy, M.D.M.; Basha, S.K.M.; Hari, M.C.; Penchalaiah, N.: Dall-e: Creating images from text. In: UGC Care Group I Journal, vol. 8, no. 14, pp. 71–75 (2021)
19. Tan, H.; Bansal, M.: Lxmert: Learning cross-modality encoder representations from transformers. arXiv preprint arXiv:1908.07490 (2019)
20. Hatamizadeh, A.; Song, J.; Liu, G.; Kautz, J.; Vahdat, A.: DiffiT: Diffusion Vision Transformers for Image Generation. arXiv preprint arXiv:2312.02139 (2023)
21. Fish, E.; Weinbren, J.; Gilbert, A.: Two-stream transformer architecture for long video understanding. arXiv preprint arXiv:2208.01753 (2022)
22. Luo, S.; Bimbo, J.; Dahiya, R.; Liu, H.: Robotic tactile perception of object properties: A review. *Mechatronics*, 48, 54–67 (2017)
23. Dahiya, R.S.; Metta, G.; Valle, M.; Sandini, G.: Tactile sensing—from humans to humanoids. *IEEE Transactions on Robotics*, 26(1), 1–20 (2009)
24. Kappasov, Z.; Corrales, J.-A.; Perdereau, V.: Touch driven controller and tactile features for physical interactions. *Robotics and Autonomous Systems*, 123, 103332 (2020)
25. Zhu, F.; Jia, R.; Yang, L.; Yan, Y.; Wang, Z.; Pan, J.; Wang, W.: Visual-tactile sensing for real-time liquid volume estimation in grasping. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 12542–12549. IEEE (2022)
26. Kelestemur, T.; Platt, R.; Padir, T.: Tactile pose estimation and policy learning for unknown object manipulation. arXiv preprint arXiv:2203.10685 (2022)
27. Xie, S.; Zhang, Y.; Zhang, H.; Jin, M.: "Development of triaxis electromagnetic tactile sensor with adjustable sensitivity and measurement range for robot manipulation". *IEEE Transactions on Instrumentation and Measurement*, 71, pp. 1–9 (2022)

28. Yuan, W.; Dong, S.; Adelson, E.H.: Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12), 2762 (2017)
29. Gomes, D.F.; Lin, Z.; Luo, S.: GelTip: A finger-shaped optical tactile sensor for robotic manipulation. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 9903–9909. IEEE (2020)
30. Ward-Cherrier, B.; Pestell, N.; Cramphorn, L.; Winstone, B.; Giannaccini, M.E.; Rossiter, J.; Lepora, N.F.: The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies. *Soft Robotics*, 5(2), 216–227 (2018)
31. Lambeta, M.; Chou, P.-W.; Tian, S.; Yang, B.; Maloon, B.; Most, V.R.; Stroud, D.; Santos, R.; Byagowi, A.; Kammerer, G.; et al.: Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3), 3838–3845 (2020)
32. Church, A.; Lloyd, J.; Lepora, N.F.; et al.: Tactile sim-to-real policy transfer via real-to-sim image translation. In: Conference on Robot Learning, pp. 1645–1654. PMLR (2022)
33. Luo, S.; Yuan, W.; Adelson, E.; Cohn, A.G.; Fuentes, R.: Vitac: Feature sharing between vision and tactile sensing for cloth texture recognition. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 2722–2727. IEEE (2018)
34. Pecyna, L.; Dong, S.; Luo, S.: Visual-tactile multimodality for following deformable linear objects using reinforcement learning. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 3987–3994. IEEE (2022)
35. She, Y.; Wang, S.; Dong, S.; Sunil, N.; Rodriguez, A.; Adelson, E.: "Cable manipulation with a tactile-reactive gripper". *The International Journal of Robotics Research*, 40(12-14), pp. 1385–1401 (2021). Publisher: SAGE Publications Sage UK: London, England.
36. Dong, S.; Jha, D.K.; Romeres, D.; Kim, S.; Nikovski, D.; Rodriguez, A.: Tactile-rl for insertion: Generalization to objects of unknown geometry. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 6437–6443. IEEE (2021)
37. Han, Y.; Batra, R.; Boyd, N.; Zhao, T.; She, Y.; Hutchinson, S.; Zhao, Y.: Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer. *arXiv preprint arXiv:2112.06374* (2021)
38. Wei, F.; Zhao, J.; Shan, C.; Yuan, Z.: Alignment and multi-scale fusion for visual-tactile object recognition. In: 2022 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE (2022)
39. Chen, Y.; Sipos, A.; Van der Merwe, M.; Fazeli, N.: Visuo-tactile transformers for manipulation. *arXiv preprint arXiv:2210.00121* (2022)
40. Khan, S.; Naseer, M.; Hayat, M.; Zamir, S.W.; Khan, F.S.; Shah, M.: Transformers in vision: A survey. *ACM Computing Surveys (CSUR)*, 54(10s), 1–41 (2022)
41. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019)
42. Wang, S.; Li, B.; Khabsa, M.; Fang, H.; Ma, H.: Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768* (2020)
43. Fedus, W.; Zoph, B.; Shazeer, N.: Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1), 5232–5270 (2022)
44. He, K.; Zhang, X.; Ren, S.; Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778 (2016)
45. Ba, J.L.; Kiros, J.R.; Hinton, G.E.: Layer normalization. *arXiv preprint arXiv:1607.06450* (2016)

46. Wang, T.; Yang, C.; Kirchner, F.; Du, P.; Sun, F.; Fang, B.: Multimodal grasp data set: A novel visual–tactile data set for robotic manipulation. *International Journal of Advanced Robotic Systems*, 16(1) (2019)
47. Gao, R.; Chang, Y.-Y.; Mall, S.; Fei-Fei, L.; Wu, J.: ObjectFolder: A Dataset of Objects with Implicit Visual, Auditory, and Tactile Representations. In: *CoRL* (2021)
48. Gao, R.; Si, Z.; Chang, Y.-Y.; Clarke, S.; Bohg, J.; Fei-Fei, L.; Yuan, W.; Wu, J.: Objectfolder 2.0: A multisensory object dataset for sim2real transfer. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10598–10608 (2022)
49. Gao, R.; Dou, Y.; Li, H.; Agarwal, T.; Bohg, J.; Li, Y.; Fei-Fei, L.; Wu, J.: The OBJECTFOLDER BENCHMARK: Multisensory Learning with Neural and Real Objects. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17276–17286 (2023)
50. Yang, F.; Ma, C.; Zhang, J.; Zhu, J.; Yuan, W.; Owens, A.: Touch and go: Learning from human-collected vision and touch. *arXiv preprint arXiv:2211.12498* (2022)
51. Awan, M.I.; Jeon, S.: Surface texture classification based on transformer network. In: *HCI*, pp. 762–764 (2023)
52. Huang, R.; Zheng, W.; Zhang, B.; Zhou, J.; Cui, Z.; Zhang, Z.: Deep learning with tactile sequences enables fruit recognition and force prediction for damage-free grasping. *Computers and Electronics in Agriculture*, 211, 107985 (2023)
53. Zhao, Y.; Jing, X.; Qian, K.; Gomes, D.F.; Luo, S.: Skill generalization of tubular object manipulation with tactile sensing and Sim2Real learning. *Robotics and Autonomous Systems*, 160, 104321 (2023)
54. Zhu, Y.: "Perceptual Vibrotactile Signal Code Based on Transformer". In: *2022 IEEE 5th Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*, vol. 5, pp. 916–922 (2022)
55. Hassen, R.; Gülecüyüz, B.; Steinbach, E.: "Pvc-slp: Perceptual vibrotactile-signal compression based on sparse linear prediction". *IEEE Transactions on Multimedia*, 23, pp. 4455–4468 (2020)
56. Strese, M.; Brudermueller, L.; Kirsch, J.; Steinbach, E.: "Haptic material analysis and classification inspired by human exploratory procedures". *IEEE Transactions on Haptics*, 13(2), pp. 404–424 (2019)
57. Sun, H.; Chen, M.; Weng, J.; Liu, Z.; Geng, G.: "Anomaly detection for in-vehicle network using CNN-LSTM with attention mechanism". *IEEE Transactions on Vehicular Technology*, 70(10), pp. 10880–10893 (2021)
58. Hassan, W.; Abdulali, A.; Abdullah, M.; Ahn, S.C.; Jeon, S.: "Towards universal haptic library: Library-based haptic texture assignment using image texture and perceptual space". *IEEE Transactions on Haptics*, 11(2), pp. 291–303 (2017)
59. Miralles, D.; Garrofé, G.; Parés, C.; González, A.; Serra, G.; Soto, A.; de Beeck, H.O.; Masson, H.L.; Sevillano, X.: "Cross-Modal Environment Self-Adaptation During Object Recognition in Artificial Cognitive Systems" (2021)
60. Wei, F.; Zhao, J.; Shan, C.; Yuan, Z.: "Alignment and Multi-Scale Fusion for Visual-Tactile Object Recognition". In: *2022 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8 (2022)
61. Strese, M.; Lee, J.-Y.; Schuwerk, C.; Han, Q.; Kim, H.-G.; Steinbach, E.: "A haptic texture database for tool-mediated texture recognition and classification". In: *2014 IEEE International Symposium on Haptic, Audio and Visual Environments and Games (HAVE) Proceedings*, pp. 118–123 (2014)

62. Strese, M.; Schuwerk, C.; Steinbach, E.: "Surface classification using acceleration signals recorded during human freehand movement". In: 2015 IEEE World Haptics Conference (WHC), pp. 214–219 (2015)
63. Li, B.; Bai, J.; Qiu, S.; Wang, H.; Guo, Y.: Vito-transformer: A visual-tactile fusion network for object recognition. *IEEE Transactions on Instrumentation and Measurement* (2023)
64. Shi, R.; Yang, S.; Chen, Y.; Wang, R.; Zhang, M.; Lu, J.; Cao, Y.: "CNN-Transformer for visual-tactile fusion applied in road recognition of autonomous vehicles". *Pattern Recognition Letters*, 166, pp. 200–208 (2023)
65. Song, R.; Sun, X.; Liu, G.: "Cross-Modal Generation of Tactile Friction Coefficient from Audio and Visual Measurements by Transformer". *IEEE Transactions on Instrumentation and Measurement* (2023)
66. Allen, J.B.: The inverse short-time Fourier transform (ISTFT) and its applications. In: 13th International Conference on Theoretical and Computational Acoustics, ICTCA 2017, Institute of Mechanics and Mechatronics, Faculty of Mechanical, TU Wien (2017)
67. Wei, X.; Shi, Y.; Zhou, L.: "Haptic signal reconstruction for cross-modal communications". *IEEE Transactions on Multimedia*, 24, pp. 4514–4525 (2021)
68. Gangisetty, S.; Rai, R.R.: Underwater image restoration using deep encoder–decoder network with symmetric skip connections. *Signal, Image and Video Processing*, pp. 1–9 (2022)
69. Cai, S.; Zhu, K.: Multi-modal transformer-based tactile signal generation for haptic texture simulation of materials in virtual and augmented reality. In: 2022 IEEE International Symposium on M
70. Li, J.; Dong, S.; Adelson, E.: "Slip detection with combined tactile and visual information". In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 7772–7777 (2018)
71. Calandra, R.; Owens, A.; Jayaraman, D.; Lin, J.; Yuan, W.; Malik, J.; Adelson, E.H.; Levine, S.: More than a feeling: Learning to grasp and regrasp using vision and touch. *IEEE Robotics and Automation Letters*, 3(4), 3300–3307 (2018)
72. Zadeh, A.; Liang, P.P.; Poria, S.; Vij, P.; Cambria, E.; Morency, L.-P.: Multi-attention recurrent network for human communication comprehension. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32, no. 1 (2018)
73. Cui, S.; Wei, J.; Li, X.; Wang, R.; Wang, Y.; Wang, S.: Generalized visual-tactile transformer network for slip detection. *IFAC-PapersOnLine*, 53(2), 9529–9534 (2020)
74. Qi, H.; Yi, B.; Suresh, S.; Lambeta, M.; Ma, Y.; Calandra, R.; Malik, J.: General in-hand object rotation with vision and touch. *arXiv preprint arXiv:2309.09979* (2023)
75. Chen, Y.; Sipos, A.; Van der Merwe, M.; Fazeli, N.: "Visuo-Tactile Transformers for Manipulation". *arXiv preprint arXiv:2210.00121* (2022)
76. Noguchi, W.; Iizuka, H.; Yamamoto, M.: "Multimodal Plastic Body and Peripersonal Space Representation Developed Through Learning of Visuo-Tactile-Proprioceptive Sensations". In: ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference (2023)
77. Li, C.: "Openai's gpt-3 language model: A technical overview". *Blog Post* (2020)
78. Li, B.; Pandey, S.; Fang, H.; Lyv, Y.; Li, J.; Chen, J.; Xie, M.; Wan, L.; Liu, H.; Ding, C.: Ftrans: energy-efficient acceleration of transformers using fpga. In: Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design, pp. 175–180 (2020)
79. Wang, H.; Wu, Z.; Liu, Z.; Cai, H.; Zhu, L.; Gan, C.; Han, S.: Hat: Hardware-aware transformers for efficient natural language processing. *arXiv preprint arXiv:2005.14187* (2020)

80. Bender, G.; Kindermans, P.-J.; Zoph, B.; Vasudevan, V.; Le, Q.: Understanding and simplifying one-shot architecture search. In: International conference on machine learning, pp. 550–559. PMLR (2018)
81. Guo, Z.; Zhang, X.; Mu, H.; Heng, W.; Liu, Z.; Wei, Y.; Sun, J.: Single path one-shot neural architecture search with uniform sampling. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16, pp. 544–560. Springer (2020)
82. Yao, Q.; Xu, J.; Tu, W.-W.; Zhu, Z.: Efficient neural architecture search via proximal iterations. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, no. 04, pp. 6664–6671 (2020)
83. Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; Carreira, J.: Perceiver: General perception with iterative attention. In: International conference on machine learning, pp. 4651–4664. PMLR (2021)
84. He, K.; Zhang, X.; Ren, S.; Sun, J.: "Deep residual learning for image recognition". In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
85. Jaegle, A.; Borgeaud, S.; Alayrac, J.-B.; Doersch, C.; Ionescu, C.; Ding, D.; Koppula, S.; Zoran, D.; Brock, A.; Shelhamer, E.; et al.: Perceiver IO: A general architecture for structured inputs & outputs. arXiv preprint arXiv:2107.14795 (2021)