

BABEL: Bodies, Action and Behavior with English Labels

Abhinanda R. Punnakal^{*,1} Arjun Chandrasekaran^{*,1} Nikos Athanasiou¹
Alejandra Quirós-Ramírez² Michael J. Black¹

¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²Universität Konstanz, Konstanz, Germany
{apunnakkal, achandrasekaran, nathanasiou, alejandra.quirós, black}@tue.mpg.de

Abstract

Understanding the semantics of human movement – the what, how and why of the movement – is an important problem that requires datasets of human actions with semantic labels. Existing datasets take one of two approaches. Large-scale video datasets contain many action labels but do not contain ground-truth 3D human motion. Alternatively, motion-capture (mocap) datasets have precise body motions but are limited to a small number of actions. To address this, we present BABEL, a large dataset with language labels describing the actions being performed in mocap sequences. BABEL consists of language labels for over 43 hours of mocap sequences from AMASS, containing over 250 unique actions. Each action label in BABEL is precisely aligned with the duration of the corresponding action in the mocap sequence. BABEL also allows overlap of multiple actions, that may each span different durations. This results in a total of over 66000 action segments. The dense annotations can be leveraged for tasks like action recognition, temporal localization, motion synthesis, etc. To demonstrate the value of BABEL as a benchmark, we evaluate the performance of models on 3D action recognition. We demonstrate that BABEL poses interesting learning challenges that are applicable to real-world scenarios, and can serve as a useful benchmark for progress in 3D action recognition. The dataset, baseline methods, and evaluation code are available and supported for academic research purposes at <https://babel.is.tue.mpg.de/>.

1. Introduction

A key goal in computer vision is to understand human movement in semantic terms. Relevant tasks include predicting semantic labels for a human movement, e.g., action recognition [12], video description [37], temporal localization [26, 39], and generating human movement that is conditioned on semantics, e.g., motion synthesis conditioned




Figure 1. People moving naturally often perform multiple actions simultaneously, and sequentially, with transitions between them. Existing large-scale 3D mocap datasets, however, describe an entire sequence with only a single action label. In BABEL, all frames and all actions are labeled. Each frame label is precisely aligned with the frames representing the action (colored wedge). This includes simultaneous actions (overlapping wedges) and transitions between actions (gray wedges).

on actions [10], or sentences [3, 17].

Large-scale datasets that capture variations in human movement and language descriptions that express the semantics of these movements, are critical to making progress on these challenging problems. Existing datasets contain detailed action descriptions for only 2D videos, e.g., ActivityNet [26], AVA [9] and HACS [39]. The large scale 3D datasets that contain action labels, e.g., NTU RGB+D 60 [27] and NTU RGB+D 120 [40] do not contain ground truth 3D human motion but only noisy estimates. On the other hand, motion-capture (mocap) datasets [2, 8, 11, 14] are small in scale and are only sparsely labeled with very few actions. We address this shortcoming with BABEL, a large dataset of diverse, densely annotated, actions with labels for all the actions in a motion capture (mocap) sequence.

We acquire action labels for sequences in BABEL, at two different levels of resolution. Similar to existing mocap datasets, we collect a sequence label that describes the action being performed in the entire sequence, e.g., `jump over obstacle` in Fig. 1. At a finer-grained resolution, the frame labels describe the action being performed at each frame of the sequence, e.g., `stand`, `run`, etc. The frame

* Denotes equal contribution.

labels are precisely aligned with the corresponding frames in the sequence that represent the action. BABEL also captures simultaneous actions, e.g., `jump over obstacle` and `support body with right hand`. We ensure that all frames in a sequence are labeled with at least one action, and all the actions in a frame are labeled. This results in dense action annotations for high-quality mocap data.

BABEL leverages the recently introduced AMASS dataset [20] for mocap sequences. AMASS is a large corpus of mocap datasets that are unified with a common representation. It has > 43 hours of mocap data (13220 sequences) performed by over 346 subjects. The scale and diversity of AMASS presents an opportunity for data-driven learning of semantic representations for 3D human movement.

Most existing large-scale datasets with action labels [2, 11, 19, 21, 26, 27, 39, 40] first determine a fixed set of actions that are of interest. Following this, actors performing these actions are captured (3D datasets), or videos containing the actions of interest are mined from the web (2D datasets). While this ensures the presence of the action of interest in the sequence, all other actions remain unlabeled. The sparse action label for a sequence, while useful, serves only as weak supervision for data-driven models that aim to correlate movements with semantic labels. This is suboptimal. 3D datasets such as NTU RGB+D [19, 27] and HumanAct12 [40] handle this shortcoming by cropping out segments that do not correspond to the action of interest from natural human movement sequences. While the action labels for the short segments are accurate, the cropped segments are unlike the natural, continuous human movements in the real-world. Thus, the pre-segmented movements are less suitable as training data for real-world applications.

Our key idea with BABEL is that natural human movement often involves multiple actions and transitions between them. Thus, understanding the semantics of natural human movement not only involves modeling the relationship between an isolated action and its corresponding movement but also the relationship between different actions that occur simultaneously and sequentially. With BABEL, our goal is to provide accurate data for statistical learning, which reflects the variety, concurrence and temporal compositions of actions in natural human movement.

The current version of BABEL contains dense action annotations for about 43.5 hours of mocap from AMASS, with 9421 unique language labels. Via a semi-automatic process of semantic clustering followed by manual categorization, we organize these into 252 actions such as `greet`, `hop`, `scratch`, `dance`, `play instrument`, etc. These actions belong to 8 broad semantic categories involving simple actions (`throw`, `jump`), complex activities (`martial arts`, `dance`), body part interactions (`scratch`, `touch face`), etc. (see Sec. 3.4). There are a total of 66289 action segments, and a single mocap sequence has 5.01 segments

on average, with 3.4 unique actions. We collect the action labels and alignments by adapting an existing web annotation tool, VIA [6] (see Sec. 3.1). Labelling was done by using Amazon Mechanical Turk [1].

We benchmark the performance of models on BABEL for the 3D action recognition task [27]. The goal is to predict the action category, given a segment of mocap that corresponds to a single action span. Unlike existing datasets containing carefully constructed segments for the actions of interest, action recognition with BABEL more closely resembles real-world applications due to the long-tailed distribution of classes in BABEL. We demonstrate that BABEL presents interesting learning challenges for an existing action recognition model that performs well on NTU RGB+D 60. In addition to being a useful benchmark for action recognition, we believe that BABEL can be leveraged by the community for tasks like pose estimation, motion synthesis, temporal localization, few shot learning, etc.

In this work, we make the following contributions: (1) We provide the largest 3D dataset of dense action labels that are precisely aligned with their corresponding movement spans in the mocap sequence. (2) We categorize the raw language labels into over 250 action classes that can be leveraged for tasks requiring categorical label sets such as 3D action recognition. (3) We analyze the actions occurring in BABEL sequences in detail, furthering our semantic understanding of mocap data that is already widely used in vision tasks. (4) We benchmark the performance of baseline 3D action recognition models on BABEL, demonstrating that the distribution of actions that resembles real-world scenarios, poses interesting learning challenges. (5) The dataset, baseline models and evaluation code are publicly available for academic research purposes at <https://babel.is.tue.mpg.de/>.

2. Related Work

Language labels and 3D mocap data. We first briefly review the action categories in large-scale 3D datasets, followed by a more detailed comparison in Table 1. The CMU Graphics Lab Motion Capture Database (CMU) [2] is widely used, and has 2605 sequences. The dataset has 6 semantic categories (e.g., ‘human interaction’, ‘interaction with environment’) that, overall, contain 23 subcategories, e.g., ‘two subjects’, ‘playground’, ‘pantomime’. Human3.6M [14] consists of 12 everyday actions in 6 semantic categories such as ‘walking variations’ (‘walking dog’, ‘walking pair’), ‘full body upright variations’ (‘greeting’, ‘posing’), etc. MoVi [8] consists of everyday actions and sports movements e.g., ‘clapping hands’, ‘pretending to take picture’, etc. KIT Whole-Body Human Motion Database (KIT) [21] focuses on both human movement and human-object interaction [32] containing grasping and manipulation actions in addition to activities such

Dataset	GT motion?	# Actions	# Hours	Per-frame?	Continuous?
CMU MoCap [2]	✓	23	9	✗	✓
MoVi [8]	✓	20	9	✗	✓
Human3.6M [14]	✓	17	18	✗	✓
LaFan1 [11]	✓	12	4.6	✗	✓
HumanAct12 [10]	✗	12	6	✗	✗
NTU RGB+D 60 [27]	✗	60	37	✗	✗
NTU RGB+D 120 [19]	✗	120	74	✗	✗
BABEL (ours)	✓	252	43.5	✓	✓

Table 1. Comparison of existing datasets containing action labels for human movement. *GT motion* indicates whether the human movements are accurate (mocap) or noisy estimates (e.g., via tracking). *# Actions* indicates the total count of action categories in each dataset. *# Hours* indicates the total duration of all sequences in the dataset. *Per-Frame?* indicates whether the action labels are precisely aligned with the corresponding spans of movement in the sequence. *Continuous?* indicates whether the movement sequences are original, continuous, human movements or short cropped segments containing specific actions. BABEL uniquely provides large-scale dense (per-frame) action labels for natural, continuous, ground-truth human movement data.

as climbing and playing sports. LaFan1 [11] is a recent dataset containing 15 different actions, including locomotion on uneven terrain, free dancing, fight movements, etc. These characterize the movement in the entire mocap sequence via simple tags or keywords. In contrast, the KIT Motion-Language Dataset [21] describes motion sequences with natural language sentences, e.g., ‘A person walks backward at a slow speed’. While our motivation to learn semantic representations of movement is similar, action labels in BABEL are precisely aligned with the sequence.

Frame actions labels in 3D mocap. The CMU MMAC dataset [30] contains precise frame labels for a fixed set of 17 cooking actions (including ‘none’). Arikan et al. [4] and Muller et al. [23] partially automate labeling temporal segments for mocap using action classifiers. While these works assume a known, fixed set of classes, in BABEL, we identify and precisely label all actions that occur in each frame.

Action labels and tracked 3D data. NTU RGB+D 60 [27] and 120 [19] are large, widely used datasets for 3D action recognition. In NTU RGB+D, RGBD sequences are captured via 3 Kinect sensors which track joint positions of the human skeleton. NTU RGB+D has segmented sequences corresponding to specific actions. There are 3 semantic categories – ‘Daily actions’ (‘drink water’, ‘taking a selfie’), ‘Medical conditions’ (‘sneeze’, ‘falling down’) and ‘Mutual actions’ (‘hugging’, ‘cheers and drink’). These datasets contain short cropped segments of actions, which differ from BABEL sequences, which are continuous, reflecting natural human movement data. The ability to model actions that can occur simultaneously, sequentially and the transitions between them is important for application to real world data [26]. See Table 1 for further comparison.

2D temporal localization. Many works over the years have contributed to progress in the action localization task [13, 29, 38]. ActivityNet [26] contains 648 hours of videos and 200 human activities that are relevant to daily life, organized under a rich semantic taxonomy. It has 19,994

(untrimmed) videos, with an average of 1.54 activities per video. More recently, HACS [39] provides a larger temporal localization dataset with 140,000 segments of actions that are cropped from 50,000 videos that span over 200 actions. AVA [9] is another recent large-scale dataset that consists of dense annotations for long video sequences for 80 atomic classes. In [36], the authors introduce a test recorded by Kinect v2 in which they describe activities as compositions of action interactions with different objects. While BABEL also contains temporally annotated labels, it does not assume a fixed set of actions that are of interest. On the other hand, with BABEL, we elicit labels for all actions in the sequence including high-level (‘eating’), and low-level actions (‘raise right hand to mouth’).

3. Dataset

We first provide details regarding the crowdsourced data collection process. We then describe the types of labels in BABEL, and the label processing procedure.

3.1. Data collection

We collect BABEL by showing rendered videos of mocap sequences from AMASS [20] to human annotators and eliciting action labels (Fig. 2). We observe that a sequence labeled as `pick up object` often also involves other actions such as walking to the object, bending down to pick up the object, grasping the object, straightening back up, turning around and walking away. We argue that labeling the entire sequence with the single label is imprecise, and problematic. First, many actions such as `turn` and `grasp` are ignored and remain unlabeled although they may be of interest to researchers [32]. Second, sequence labels provide weak supervision to statistical models, which are trained to map the concept of `picking up object` to the whole sequence when it, in fact, contains many different actions. To illustrate this point, we examine a typical sequence (see the Sup. Mat.), and find that only 20% of the du-

ration of the sequence labeled as pick up and place object corresponds to this action. Crucially, walking towards and away from the object – actions that remain unlabeled – account for 40% of the duration. While this makes semantic sense to a human – picking up and placing an object is the only action that changes the state of the world and hence worth mentioning, this might be suboptimal training data to a statistical model, especially when the dataset also contains the confusing classes walk, turn, etc. Finally, using noisy labels as ground truth during evaluation does not accurately reflect the capabilities of models.

We address this with action labels at two levels of resolution – a label describing the overall action in the entire sequence, and dense action labels that are aligned with their corresponding spans of movement in the mocap sequence.

3.2. BABEL action labels

We collect BABEL labels in a two-stage process – first, we collect sequence labels, and determine whether the sequence contains multiple actions. We then collect frame labels for the sequences containing multiple actions.

Sequence labels. In this labeling task, annotators answer two questions regarding a sequence. We first ask annotators if the video contains more than one action (yes/no).¹ If the annotator chooses ‘no’, we ask them to name the action in the video. If they instead choose ‘yes’, we elicit a sequence label with the question, “If you had to describe the whole sequence as one action, what would it be?” Please see the web-based task interface provided in the Sup. Mat.

We ask annotators to enter the sequence labels in a text-box, with the option of choosing from an auto-complete drop-down menu that is populated with a list of basic actions. We specifically elicit free-form labels (as opposed to a fixed list of categories) from annotators to discover the diversity in actions in the mocap sequences. We find that in most cases, annotators tend to enter their own action labels. This also presents a challenge, acting as a source of label variance. Apart from varying vocabulary, free-form descriptions are subject to ambiguity regarding the ‘correct’ level in the hierarchy of actions [9], e.g., raise left leg, step, walk, walk backwards, walk backwards stylishly, etc.

Overall, 2 annotators label each sequence, resulting in $13220 * 2 = 16440$ sequence labels. For 47.7% of the sequences, both annotators agree that they contain multiple actions. These are further annotated with frame labels.

Frame labels. We obtain labels for all actions that occur in the sequence, and precisely identify the span in the sequence that corresponds to the action. We leverage an existing video annotation tool, VIA [6], and modify the front-end interface and back-end functionality to suit our anno-

¹Note that the initial ‘T-pose’ for calibration, followed by standing are considered separate actions with a transition between them.

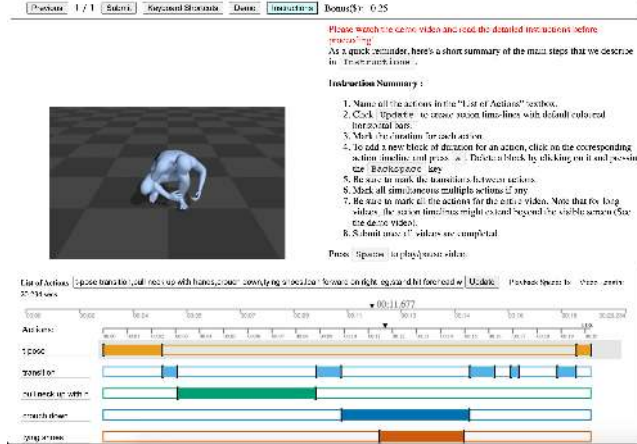


Figure 2. BABEL annotation interface to collect frame-level action labels. Annotators first name all the actions in the video. They then, precisely align the length of the action segment (colored horizontal bar) with the corresponding duration of the action in the video. This provides dense action labels for the entire sequence.

tation purposes. For instance, we ensure that every frame in the sequence is annotated with at least one action label. This includes ‘transition’ which indicates a transition between two actions, or ‘unknown’ which indicates that the annotator is unclear as to what action is being performed. This provides us with dense annotations of action labels for the sequence. Figure 2 shows a screenshot of the AMT task interface for frame label annotation in BABEL.

To provide frame labels, an annotator first watches the whole video and enters all the actions in the ‘List of Actions’ text-box below the video. This populates a set of empty colored box outlines corresponding to each action. The annotator then labels the span of an action by creating a segment (colored rectangular box) with a button press. The duration of the segment and the start/end times can be changed via simple click-and-drag operations. The video frame is continuously updated to the appropriate time-stamp corresponding to the end time of the current active segment. This provides the annotator real-time feedback regarding the exact starting point of the action. Once the segment is placed, its precision can be verified by a ‘play segment’ option that plays the video span corresponding to the current segment. In case of errors, the segment can be further adjusted. We provide detailed instructions via text, and a video tutorial that explains the task with examples, and demonstrates operation of the annotation interface. The web interface of the task is provided in the Sup. Mat.

Overall, BABEL contains labels for a total of 66289 segments² that correspond to the 13220 sequences in the dataset, with an average of 5.01 segments per sequence.

²Note that this includes labels collected via the ‘sequence label’ task. If an entire sequence has only a single action, it counts as 1 segment.

3.3. Annotators

We recruit all annotators for our tasks via the Amazon Mechanical Turk (AMT)³ crowd-sourcing platform. These annotators are located either in the US or Canada. In the sequence label annotation task, we recruit 571 unique annotators with > 5000 HITs approved and an approval rate > 95%. In the frame labeling task, which is more involved, we first run small-scale tasks to recruit annotators. For further tasks, we only qualify 133 annotators who demonstrate an understanding of the task and provide satisfactory action labels and precise segments in the sequence. In both tasks, our median hourly pay is \sim \$12. We also provide bonus pay as an incentive for thorough work in the frame labeling task (details in Sup. Mat.).

3.4. Label processing

BABEL contains a total of 9421 unique raw action labels. Note that while each action label is a unique string, labels are often semantically similar (*walk*, *stroll*, etc.), are minor variations of an action word (*walking*, *walked*, etc.) or are misspelled. Further, tasks like classification require a smaller categorical label set. We organize the raw labels into two smaller sets of semantically higher-level labels – action categories, and semantic categories.

Action categories. We map the variants of an action into a single category via a semi-automatic process that involves clustering the raw labels, followed by manual adjustment.

We first pre-process the raw string labels by lower-casing, removing the beginning and ending white-spaces, and lemmatization. We then obtain semantic representations for the raw labels by projecting them into a 300D space via Word2Vec embeddings [22]. Word2Vec is a widely used word embedding model that is based on the distributional hypothesis – words with similar meanings have similar contexts. Given a word, the model is trained to predict surrounding words (context). An intermediate representation from the model serves as a word embedding for the given word. For labels with multiple words, the overall representation is the mean of the Word2Vec embeddings of all words in the label. Labels containing words that are semantically similar, are close in the representation space.

We cluster labels that are similar in the representation space via K-means ($K = 200$ clusters). This results in several semantically meaningful clusters, e.g., *walk*, *stroll*, *stride*, etc. which are all mapped to the same cluster. We then manually verify the cluster assignments and fix them to create a semantically meaningful organization of the action labels. Raw labels that are not represented by Word2Vec (e.g., *T-pose*) are manually organized into relevant categories in this stage. For each cluster, we determine a category name that is either a syn-

onym (*walk* \leftarrow {*walk*, *stroll*, *stride*}) or hypernym (*walk* \leftarrow {*walk forward*, *walk around*}) that describes all action labels in the cluster.

Some raw labels, e.g., *rotate wrists* can be composed into multiple actions like *circular movement* and *wrist movement*. Thus, raw labels are occasionally assigned membership to multiple action categories.

Overall, the current version of BABEL has 252 action categories. Interestingly, the most frequent action in BABEL is ‘transition’ – a movement that usually remains unlabeled in most datasets. There are 17287 transitions between different actions in BABEL. Unsurprisingly, the frequency of actions decreases exponentially following Zipf’s law – the 50th most frequent action *sports moves* occurs 280 times, the 100th most frequent action *communicate*, occurs 52 times, and the 200th most frequent action *lick*, occurs 5 times. We visualize the action categories containing the largest number of raw labels (cluster elements) in Fig. 3 (outer circle). Raw labels corresponding to these categories are shown on the right. We provide histograms of duration and number of sequences per-action, in the Sup. Mat.

Semantic categories of labels. Action categories often reflect qualitatively different types of actions like interacting with objects, actions that describe the trajectory of movement, complex activities involving multiple actions, etc. We formalize the different types of actions in BABEL into 8 semantic categories (inner circle in Fig. 3):

1. **Simple dynamic actions** contain low-level atomic actions – *walk*, *run*, *kick*, *punch*, etc.
2. **Static actions** involve transitioning to and/or maintaining a certain posture – *sit*, *stand*, *kneel*, etc.
3. **Object interaction:** e.g., *place something*, *move something*, *use object*, etc.
4. **Body part interaction** contains actions like *touching face*, *scratch*, etc. which typically involve self-contact of body parts.
5. **Body part** describes the movement of a specific body part – *raise arm*, *lower head*, *rotate wrist*, etc.
6. **Type of movement** contains actions that describe the trajectory of movement of either a body part or the whole body – *twist*, *circular movement*, etc.
7. **Activity** contains complex actions that often involve multiple low-level actions – *play sports*={*run*, *jump*}, *dance*={*stretch*, *bend*}, etc.
8. **Abstract actions** contain actions which often refer to the emotional state of the person and whose physical realizations could have large variance – *excite*, *endure*, *learn*, *find*, etc. There are only a few abstract actions in BABEL.

The diversity in the types of action labels in BABEL can be leveraged by tasks modeling movement at various levels of semantic abstraction, e.g., movement of body parts

³<https://www.mturk.com/>




Figure 3. *Left.* 2D t-SNE [34] visualization of the semantic space that we project raw labels into. Similar labels are grouped via K-means clustering. The green shading and points represent the ‘dance’ cluster and its members respectively. *Center.* Distribution of (a subset of) action categories (outer circle) under each semantic category (inner circle) in BABEL. The angle occupied by the action category is proportional to the number of unique raw label strings associated with it. Action categories with a large number of fine-grained descriptions are shown. *Right.* Subset of the fine-grained descriptions associated with selected action categories.

like circular movement of wrist at a low level, or high-level semantic activities such as dancing the waltz. Further, depending on the task and model, one can exploit either the discrete set of action categories (e.g., action recognition), or embed the raw action labels into a semantic space to provide a semantic representation of the segment of movement (e.g., action synthesis).

We provide the full set of semantic categories, action categories, and raw action labels in BABEL in the Sup. Mat.

4. Analysis

Natural human movement often contains multiple actions and transitions between them. Modeling the likelihood of simultaneous actions and action transitions has applications in reasoning about action affordances in robotics and virtual avatars, motion synthesis [33], activity forecasting [16], animation [31], and action recognition.

4.1. Simultaneous actions

Although people often perform multiple actions simultaneously in real life, this is rarely captured in labeled datasets. Recall from Sec. 3.2 that in BABEL, we ask annotators to label all actions that are occurring in each frame of the sequence. Overall, BABEL has 3208 simultaneous actions, which are defined as actions that overlap for a duration of > 0.1 seconds. We exclude the overlap of an action with transition since this implies adjacent actions.

Simultaneous actions often exhibit relationships such as:

1. **Hierarchical.** Some simultaneous actions reflect the hierarchical structure in actions. For instance, a complex activity & action comprising the activity, e.g., eating food & raise right hand to

mouth, and dancing & extend arms.

2. **Complementary.** The two actions are independent, e.g., hold with left hand & look right.
3. **Superimposed.** An action can move a certain body part that partly modifies another action; e.g., carry with right hand modifies the complex activity walk, and right high kick partly modifies the static (full body) action fight stance.
4. **Compositional.** Actions involving the same body parts that result in a body or part movement that is a function of both actions, e.g., walking & turn.

4.2. Temporally adjacent actions

The dense labels in BABEL capture the progression of actions in mocap sequences. We analyze adjacent actions where where action a_i follows a_j (denoted by $a_j \rightarrow a_i$). a_i and a_j denote action segments, i.e., a contiguous set of frames corresponding to an action (and not the action for a single frame). Thus, $a_i \neq a_j$ if the actions are adjacent. We say $a_j \rightarrow a_i$ if the frame succeeding the last frame of a_j is the first frame of a_i . In practice, we account for imprecise human temporal annotations by ignoring a small overlap in duration (< 0.1 sec.) between a_i and a_j . We also disregard the separation of actions by transition; i.e., $a_j \rightarrow a_i$ if $a_j \rightarrow a_t$ and $a_t \rightarrow a_i$, where $a_t = \text{transition}$.

We visualize the frequent transitions between actions, i.e., $a_j \rightarrow a_i$ sorted by $\text{Count}(a_j \rightarrow a_i)$ in BABEL, in Fig. 4. We observe that walk, unsurprisingly, has the most diverse set of adjacent actions, i.e., $\text{Count}(\text{walk} \rightarrow a_i)$ and $\text{Count}(a_j \rightarrow \text{walk})$ are large. While transitions between action pairs such as (jog, turn), (walk, t-pose) are bidirectional (with \sim equal frequency), others have




Figure 4. Node represent actions, and an edge represents a transition between these actions in the mocap sequence. Edge thickness $\propto \text{Count}(a_i \rightarrow a_j)$ (frequency of transition) in BABEL.

#	Transition of actions
1	walk, transition, pick up, set down, transition, walk clockwise, transition, stand
2	a-pose, transition, wave hands in and out, wave arms in front of left, transition, cross left leg in circle gesture series, transition, t-pose
3	looking left, standing, transition, looking right, standing
4	stepping forward, standing, turning back, walking back, walking forward, standing, losing balance, transition, turning around, walking, standing
5	step back, stand, transition, walk to the left

Table 2. Random walk samples based on action transition probabilities learned from BABEL. The generated samples are plausible action sequences simulating natural human movement.

fewer adjacent actions. Some action categories with few transitions illustrate semantically meaningful action chains, e.g., *sit* \rightarrow *stand up* \rightarrow *walk* and *walk* \rightarrow *bend* \rightarrow *pick something up* \rightarrow *place something*. Unidirectional transitions such as *sit* \rightarrow *stand up* and *walk* \rightarrow *sit* implicitly indicate the arrow of time [24]. Interestingly, the transition from *sit* \rightarrow *stand up*, and the lack of transition from *sit* \rightarrow *stand* delineates the subtle difference between the labels *stand* (static action of ‘maintaining an upright position’) and *stand up* (dynamic action of ‘rising into an upright position’).

Given the temporally adjacent actions in BABEL, we attempt to model the transition probabilities between actions, i.e., $P(a_i|a_j)$. Concretely, we compute $P(a_i|a_{i-1}, a_{i-2}, a_{i-3})$, an order 3 Markov Chain [7], and observe in Table 2 that random walks along this chain generate plausible action sequences for human movement.

4.3. Bias

There are a few potential sources of bias in BABEL, which we report in the Sup. Mat. Specifically, we discuss potential biases introduced by the interface design, pay structure, and label processing method. We also analyze the

inter-annotator variation in BABEL labels by collecting 5 unique annotations for each of 29 sequences. We find that for the same sequence, annotators vary in the labeled action categories, the number of actions, and segments. In general, the variance appears to be larger for sequences of longer duration. We provide further details in the Sup. Mat.

5. Experiments

The dense action labels in BABEL can be leveraged for multiple vision tasks like pose estimation, motion synthesis, temporal localization, etc. In this section, we demonstrate the value of BABEL for the 3D action recognition task [19, 27], where the goal is to predict a single action category $y \in \mathcal{Y}$, for a given motion segment $(\mathbf{x}_t, \dots, \mathbf{x}_{t'})$.

Motion representation. A mocap sequence in AMASS, is an array of poses over time, $\mathbf{M} = (\mathbf{p}_1, \dots, \mathbf{p}_L)$, where \mathbf{p}_i are pose parameters of the SMPL-H body model [25]. For consistency with prior work, we predict the 25-joint skeleton used in NTU RGB+D [27] from the vertices of the SMPL-H mesh; see Sup. Mat. Thus, we represent a movement sequence as $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_L)$, where $\mathbf{x}_i \in \mathbb{R}^{J \times 3}$ represents the position of the $J (= 25)$ joints in the skeleton, in Cartesian co-ordinates, (x, y, z) .

Labels. In BABEL, a raw action label is mapped to the segment of human movement $\mathbf{X}_s = (\mathbf{x}_{ts}, \dots, \mathbf{x}_{te})$ corresponding to the action. Recall that a raw action label for a segment can belong to multiple action categories \mathcal{Y}_s (e.g., *rotate wrists* \rightarrow *circular movement*, *wrist movement*). Overall, BABEL contains N movement segments, and their action categories $(\mathbf{X}_s, \mathcal{Y}_s)^N$.

Architecture. We benchmark performance on BABEL with the 2-Stream-Adaptive Graph Convolutional Network (2s-AGCN) [28], a popular architecture that performs graph convolutions spatially (along bones in the skeleton), and temporally (joints across time). Crucially, the graph structure follows the kinematic chain of the skeleton in the first layer but is adaptive – the topology is a function of the layer and the sample. The model achieves good performance on both 2D and 3D action recognition. GCNs remain the architecture of choice even in more recent state-of-the-art approaches [5].

A 2s-AGCN consists of two streams with the same architecture – one which accepts joint positions, and the other, bone lengths and orientations, as input respectively. The final prediction is the average score from the two streams. On NTU RGB+D, 2s-AGCN achieves an accuracy of 88.5% on the cross-subject task. In our experiments, we use only the joint stream, which achieves 2% lower accuracy compared to 2s-AGCN [28].

Data pre-processing. We normalize the input skeleton by transforming the coordinates such that the joint position of the middle spine is the origin, the shoulder blades are parallel to the X-axis and the spine to the Y-axis, similar

# actions	Loss type	Top-1	Top-5	Top-1-norm
60	CE	44.87	71.78	17.16
	Focal	37.52	71.22	25.73
120	CE	43.55	69.54	11.32
	Focal	31.72	59.29	19.32
150	CE	42.08	67.29	9.61
	Focal	27.79	52.57	14.50

Table 3. 3D action recognition performance on different subsets of BABEL with 2s-AGCN [28]. CE indicates Cross-Entropy loss and Focal indicates Focal loss [18].

to Shahroudy et al. [27]. We follow the 2s-AGCN pre-processing approach and divide a segment \mathbf{X}_s into contiguous, non-overlapping 5 sec. chunks, \mathbf{X}_s^i , at 30fps, i.e., $\mathbf{X}_s = (\mathbf{X}_s^1, \dots, \mathbf{X}_s^K)$. Note that the number of chunks per segment, $K = \lceil \frac{te}{5 \times 30} \rceil$. If the K th chunk \mathbf{X}_s^K has duration < 5 sec., we repeat \mathbf{X}_s^K , and truncate at 5 sec.

Thus, a single sample in action recognition is a 5 sec. motion chunk $\mathbf{X}_s^i \in \mathbf{X}_s$, and the action category labeled for its corresponding segment $y \in \mathcal{Y}_s$ in BABEL.

BABEL action recognition splits. BABEL contains 252 action categories with a long-tailed distribution of samples per class, unlike other popular 3D action recognition datasets NTU RGB+D [19, 27]. To understand the challenge posed by the long-tailed distribution of action categories in BABEL, we perform experiments with 3 different datasets containing 60, 120 and 150 action categories (see Table 3). While BABEL-60 is already long-tailed, BABEL-150 is contains both extremely frequent and extremely rare classes. We randomly split the 13220 sequences in BABEL into train (60%), val. (20%) and test (20%) sets. We choose the model with best performance on the val. set, and report performance on the test set. We provide the precise distribution of action categories for each split in the Sup. Mat.

Metrics. Top-1 measures the accuracy of the highest-scoring prediction. Top-5 evaluates whether the ground-truth category is present among the top 5 highest-scoring predictions. It accounts for labeling noise and inherent label ambiguity. Note that it also accounts for the possible presence of multiple action categories \mathcal{Y}_s , per input movement sequence. Ideal models will score all the categories relevant to a sample highly. Top-1-norm is the mean Top-1 across categories. The magnitude of $(\text{Top-1-norm}) - (\text{Top-1})$ illustrates the class-specific bias in the model performance. In BABEL, it reflects the impact of class imbalance on learning.

Training. We experiment with two losses – standard cross-entropy loss, and the recently introduced focal loss [18] which compensates for class imbalance by weighting the cross-entropy loss higher for inaccurate predictions. We use the Adam optimizer [15] with a learning rate of 0.001. We train models for ~ 250 epochs with batch size = 64.

Results. We observe that the decrease in Top-1 and Top-5 performance with the increase in number of classes is relatively small, in Table 3. Importantly, we note that Top-1-norm is much lower than Top-1 . This clearly points to inefficient learning from the long-tailed class distribution in BABEL. Training with Focal significantly improves Top-1-norm performance on all BABEL subsets. This is encouraging for efforts to learn models with lower class-specific biases despite severe class imbalance.

BABEL as a recognition benchmark. On the widely used NTU RGB+D benchmark, Top-1 recognition performance approaches 87% with 2s-AGCN [28]. Note that unlike NTU RGB+D whose distribution of motions and actions is carefully controlled, the diversity and long-tailed distribution of samples in BABEL makes the recognition task more challenging. The few-shot recognition split of NTU RGB+D 120 partly addresses this issue. However, considering few-shot learning as a separate task typically involves measuring performance on only the few-shot classes, ignoring the larger distribution. Models in the real world ideally need to learn and perform well on both the frequent and infrequent classes of an imbalanced distribution [35]. We present BABEL as an additional benchmark for 3D action recognition, which evaluates the ability of models to learn from more realistic distributions of actions.

6. Conclusion

We presented BABEL, a large-scale dataset with dense action labels for mocap sequences. Unlike existing 3D datasets with action labels, BABEL has labels for *all* actions that occur in the sequence including simultaneously occurring actions, and transitions between actions. We analyzed the relationships between temporally adjacent actions and simultaneous actions occurring in a sequence. We demonstrated that the action recognition task on BABEL is challenging due to the label diversity and long-tailed distribution of samples. We believe that BABEL will serve as a useful additional benchmark for action recognition since it evaluates the ability to model realistic distributions of data. We hope that this large-scale, high quality dataset will accelerate progress in the challenging problem of understanding human movement in semantic terms.

Acknowledgements. We thank Taylor McConnell, Leyre Sánchez Vinuela, Mila Gorecki for help with categorization of labels, Muhammed Kocabaş for the interesting discussions, Cornelia Köhler and Omid Taheri for providing feedback on the manuscript. Nikos Athanasiou acknowledges funding by Max Planck Graduate Center for Computer and Information Science Doctoral Program.

Disclosure: MJB has received research gift funds from Adobe, Intel, Nvidia, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his research was performed solely at, and funded solely by, Max Planck. MJB has financial interests in Amazon, Datagen Technologies, and Meshcapade GmbH.

References

- [1] Amazon Mechanical Turk. (Date last accessed 16-November-2020). [2](#)
- [2] CMU Graphics Lab. (Date last accessed 13-November-2020). [1](#), [2](#), [3](#)
- [3] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision, 3DV 2019, Québec City, QC, Canada, September 16-19, 2019*, pages 719–728. IEEE, 2019. [1](#)
- [4] Okan Arikian, David A. Forsyth, and James F. O’Brien. Motion synthesis from annotations. *ACM Transactions on Graphics*, 22(3):402–408, 2003. [3](#)
- [5] Ke Cheng, Yifan Zhang, Xiangyu He, Weihang Chen, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with shift graph convolutional network. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 180–189. IEEE, 2020. [7](#)
- [6] Abhishek Dutta and Andrew Zisserman. The VIA annotation software for images, audio and video. In *Proceedings of the 27th ACM International Conference on Multimedia, MM 2019, Nice, France, October 21-25, 2019*, pages 2276–2279. ACM, 2019. [2](#), [4](#)
- [7] Paul A Gagniuć. *Markov chains: from theory to implementation and experimentation*. John Wiley & Sons, 2017. [7](#)
- [8] Saeed Ghorbani, Kimia Mahdavian, Anne Thaler, Konrad Kording, Douglas James Cook, Gunnar Blohm, and Nikolaus F. Troje. Movi: A large multipurpose motion and video dataset, 2020. [1](#), [2](#), [3](#)
- [9] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6047–6056. IEEE Computer Society, 2018. [1](#), [3](#), [4](#)
- [10] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2Motion: Conditioned Generation of 3d Human Motions. In *MM ’20: The 28th ACM International Conference on Multimedia, Virtual Event / Seattle, WA, USA, October 12-16, 2020*, pages 2021–2029. ACM, 2020. [1](#), [3](#)
- [11] Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher J. Pal. Robust motion in-betweening. *ACM Trans. Graph.*, 39(4):60, 2020. [1](#), [2](#), [3](#)
- [12] Samitha Herath, Mehrtash Tafazzoli Harandi, and Fatih Porikli. Going deeper into action recognition: A survey. *Image Vision Computing*, 60:4–21, 2017. [1](#)
- [13] Haroon Idrees, Amir Roshan Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The THUMOS challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017. [3](#)
- [14] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. [1](#), [2](#), [3](#)
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [8](#)
- [16] Kris M. Kitani, Brian D. Ziebart, James Andrew Bagnell, and Martial Hebert. Activity forecasting. In *Computer Vision - ECCV 2012 - 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part IV*, volume 7575 of *Lecture Notes in Computer Science*, pages 201–214. Springer, 2012. [6](#)
- [17] Angela S Lin, Lemeng Wu, Rodolfo Corona, Kevin Tai, Qixing Huang, and Raymond J Mooney. Generating animated videos of human activities from natural language descriptions. *Learning*, 2018, 2018. [1](#)
- [18] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):318–327, 2020. [8](#)
- [19] Jun Liu, Amir Shahrudiy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C. Kot. NTU RGB+D 120: A large-scale benchmark for 3d human activity understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(10):2684–2701, 2020. [2](#), [3](#), [7](#), [8](#)
- [20] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5441–5450. IEEE, 2019. [2](#), [3](#)
- [21] Christian Mandery, Ömer Terlemeç, Martin Do, Nikolaus Vahrenkamp, and Tamim Asfour. The KIT whole-body human motion database. In *International Conference on Advanced Robotics, ICAR 2015, Istanbul, Turkey, July 27-31, 2015*, pages 329–336. IEEE, 2015. [2](#), [3](#)
- [22] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119, 2013. [5](#)
- [23] Meinard Müller, Andreas Baak, and Hans-Peter Seidel. Efficient and robust annotation of motion capture data. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA 2009, New Orleans, Louisiana, USA, August 1-2, 2009*, pages 17–26. ACM, 2009. [3](#)
- [24] Lyndsey C. Pickup, Zheng Pan, Donglai Wei, Yi-Chang Shih, Changshui Zhang, Andrew Zisserman, Bernhard Schölkopf, and William T. Freeman. Seeing the arrow of time. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June*

- 23-28, 2014, pages 2043–2050. IEEE Computer Society, 2014. [7](#)
- [25] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), Nov. 2017. [7](#)
- [26] Jan Sedmidubský, Petr Elias, and Pavel Zezula. Benchmarking search and annotation in continuous human skeleton sequences. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval, ICMR 2019, Ottawa, ON, Canada, June 10-13, 2019*, pages 38–42. ACM, 2019. [1](#), [2](#), [3](#)
- [27] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3d human activity analysis. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 1010–1019. IEEE Computer Society, 2016. [1](#), [2](#), [3](#), [7](#), [8](#)
- [28] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12026–12035. Computer Vision Foundation / IEEE, 2019. [7](#), [8](#)
- [29] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, volume 9905 of *Lecture Notes in Computer Science*, pages 510–526. Springer, 2016. [3](#)
- [30] Ekaterina H. Spriggs, Fernando De la Torre, and Martial Hebert. Temporal segmentation and activity classification from first-person sensing. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2009, Miami, FL, USA, 20-25 June, 2009*, pages 17–24. IEEE Computer Society, 2009. [3](#)
- [31] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Transactions on Graphics*, 38(6):209:1–209:14, 2019. [6](#)
- [32] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part IV*, volume 12349 of *Lecture Notes in Computer Science*, pages 581–600. Springer, 2020. [2](#), [3](#)
- [33] Julian Tanke, A. Weber, and Juergen Gall. Human motion anticipation with symbolic label. *ArXiv*, abs/1912.06079, 2019. [6](#)
- [34] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. [6](#)
- [35] Davis Wertheimer and Bharath Hariharan. Few-shot learning with localization in realistic settings. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6558–6567. Computer Vision Foundation / IEEE, 2019. [8](#)
- [36] Chenxia Wu, Jiemi Zhang, Silvio Savarese, and Ashutosh Saxena. Watch-n-patch: Unsupervised understanding of actions and relations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4362–4370. IEEE Computer Society, 2015. [3](#)
- [37] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A large video description dataset for bridging video and language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5288–5296. IEEE Computer Society, 2016. [1](#)
- [38] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *Int. J. Comput. Vis.*, 126(2-4):375–389, 2018. [3](#)
- [39] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. HACS: human action clips and segments dataset for recognition and temporal localization. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 8667–8677. IEEE, 2019. [1](#), [2](#), [3](#)
- [40] Shihao Zou, Xinxin Zuo, Yiming Qian, Sen Wang, Chi Xu, Minglun Gong, and Li Cheng. Polarization human shape and pose dataset. *CoRR*, abs/2004.14899, 2020. [1](#), [2](#)