

# Cross-modal Recurrent Models for Weight Objective Prediction from Multimodal Time-series Data

**Petar Veličković**  
University of Cambridge  
and Nokia Bell Labs  
Cambridge, UK  
petar.velickovic@cst.cam.ac.uk

**Laurynas Karazija**  
University of Cambridge  
Cambridge, UK  
lk378@cam.ac.uk

**Nicholas D. Lane**  
University of Oxford  
and Nokia Bell Labs  
Oxford, UK  
nicholas.lane@cs.ox.ac.uk

**Sourav Bhattacharya**  
Nokia Bell Labs  
Cambridge, UK  
sourav.bhattacharya@  
nokia-bell-labs.com

**Edgar Liberis**  
University of Cambridge  
Cambridge, UK  
el398@cam.ac.uk

**Pietro Liò**  
University of Cambridge  
Cambridge, UK  
pietro.liò@cst.cam.ac.uk

**Angela Chieh**  
Nokia Digital Health - Withings  
Issy-les-Moulineaux, France  
angela.chieh@withings.com

**Otmane Bellahsen**  
Nokia Digital Health - Withings  
Issy-les-Moulineaux, France  
otmane.bellahsen@withings.com

**Matthieu Vegreville**  
Nokia Digital Health - Withings  
Issy-les-Moulineaux, France  
matthieu.vegreville@withings.com

## ABSTRACT

We analyse multimodal time-series data corresponding to weight, sleep and steps measurements. We focus on predicting whether a user will successfully achieve his/her weight objective. For this, we design several deep long short-term memory (LSTM) architectures, including a novel *cross-modal LSTM (X-LSTM)*, and demonstrate their superiority over baseline approaches. The X-LSTM improves parameter efficiency by processing each modality separately and allowing for information flow between them by way of *recurrent cross-connections*. We present a general hyperparameter optimisation technique for X-LSTMs, which allows us to significantly improve on the LSTM and a prior state-of-the-art cross-modal approach, using a comparable number of parameters. Finally, we visualise the model's predictions, revealing implications about latent variables in this task.

## CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in ubiquitous and mobile computing**; • **Computing methodologies** → **Neural networks**; *Supervised learning by classification*; *Cross-validation*; • **Applied computing** → **Health informatics**;

## KEYWORDS

Health data analytics, Multimodal time-series prediction, Long short-term memory, X-LSTM, Fitness applications.

## ACM Reference Format:

Petar Veličković, Laurynas Karazija, Nicholas D. Lane, Sourav Bhattacharya, Edgar Liberis, Pietro Liò, Angela Chieh, Otmane Bellahsen, and Matthieu Vegreville. 2018. Cross-modal Recurrent Models for Weight Objective Prediction from Multimodal Time-series Data. In *PervasiveHealth '18: 12th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '18)*, May 21–24, 2018, New York, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3240925.3240937>

## 1 INTRODUCTION

Recently, consumer-grade health devices, such as wearables and smart home appliances became more widespread, which presents new data modelling opportunities. Here, we investigate one such task—predicting the users' future body weight in relation to their weight goal given historical weight, and sleep and steps measurements. This study is enabled by a first-of-its-kind dataset of fitness measurements from ~15000 users. Data are captured from different sources, such as smart-watches, wrist- and hip-mounted wearables, smartphone applications and smart bathroom scales. This is one of the

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*PervasiveHealth '18*, May 21–24, 2018, New York, NY, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6450-8/18/05...\$15.00

<https://doi.org/10.1145/3240925.3240937>

first times that such quantities of large-scale longitudinal (spanning *up to 500 consecutive days* of comprehensive measurements recorded per user) multi-device consumer-grade health data have been investigated. From this dataset, we study a binary classification task: will the users be able to achieve the weight goal they input into their smart device?

There is a range of potential scenarios where such predictive modelling would be useful for weight control within consumer health systems. Motivating examples include: *direct feedback* to the user about his/her progress, *suggesting new, realistic weight objectives*, and *evaluating the effects of major lifestyle changes*. Significant work already exists towards developing consumer systems of this type [13–15, 28] but they often assume the availability of scalable predictive models of user behaviour, such as the weight goal prediction task we investigate.

Initially, the very tractability of this problem may be questioned: modelling even thousands of users limits the kind of data that sufficiently many user devices can accurately measure (such as daily step counts, or hours slept). Therefore, many factors key in weight change (such as eating habits) must remain as only latently observed. In this work, we show that that deep long short-term memory (LSTM) [7] models *are* able to produce accurate predictions in this setting, significantly outperforming baseline approaches, even though some factors are only observed latently. We also discover interesting patterns in input sequences that push the network’s confidence in success or failure to extremes. We hypothesise that these patterns affect latent variables and link our hypotheses to existing research on sleep.

Most importantly, we improve the parameter efficiency of LSTM models for multimodal input (in this case sleep, step and weight measurements) by proposing *cross-modal LSTMs (X-LSTMs)*. X-LSTMs extract features from each modality separately, while still allowing for *information flow* between the different modalities by way of *cross-connections*. Our findings are supported by a general data-driven methodology (applicable to *arbitrary* multimodal problems) that exploits unimodal predictive power to vastly simplify finding appropriate hyperparameters for X-LSTMs (reducing most of the tuning effort to a *single parameter*). We also compare our model to a previous state-of-the-art cross-modal sequential data technique [21], outlining its limitations and successfully outperforming it on this task.

This manuscript further builds up on our (non-archival) workshop paper of the same name [26], incorporating a significantly expanded discussion of methodology (regarding the motivations for this work, the description of the dataset and preprocessing, as well as the clarity of exposition of the utilised machine learning models). It also provides several additional analyses of results—quantitative (ROC curves, impact of micro-architectural decisions) as well as qualitative

(analysing the effect of weight objective magnitudes to the predictive biases of the model).

## 2 DATASET AND PREPROCESSING

We performed our investigation on anonymised data obtained from several devices across the *Nokia Digital Health - Withings* range. The dataset contains weight, height, sleep and steps measurements, as well as user specified weight objectives. Weights are measured by the Withings scale. All other data are obtained from the Withings application through the use of wearables.

Users were first included in the dataset under the condition of having recorded at least 10 weight measurements over a 2-month period. In total, the dataset contains 1 664 877 such users. Further processing was performed to remove outliers or those users with too few, or too sporadic, data observations; after this stage  $\sim 15\text{K}$  users were remaining. The precise steps taken to reach this final dataset are described below.

Obvious outliers, reports of unrealistic heights (below 130 cm or above 225cm), and/or consistent weight changes of more than 1.5kg per day have been discarded. Steps and sleep are recorded on a per-day basis, while weights are recorded at the user’s discretion; to align the weight measurements with the other two modalities, we have applied a moving average to the person’s recorded weight throughout an individual day. A sequence may be labelled with any weight objective that has been set by the user, and is still unachieved, by the time the sequence ends. Overly ambitious objectives (over  $\pm 20$  kilograms proposed) are ignored.

We consider a weight objective *successful* if there exists a weight measurement in the future that reaches or exceeds it, and we consider it *unsuccessful* if the user *stops recording weights* (allowing for a long enough window after the end of the recorded sequence) or *sets a more conservative objective* in the meantime. In line with known best practices in deep learning, data are normalised to have mean zero and standard deviation one per-feature.

The derived dataset spans 18036 sequences associated with weight objectives. All of the sequences are comprised of user-related features: height, gender, age category, weight objective; along with sequential features—for each day: duration of light and deep sleep, time to fall asleep and time spent awake; number of times awoken during the night; time required to wake up; bed-in/bed-out times; steps and (average) weights for the day. We consider sequences that span at least 10 contiguous days. There are  $\sim 3.6\text{K}$  users for whom more than one sequence is present in the dataset.

Every sequence also has a boolean *label*, indicating whether the objective has been successfully achieved at some point in the future. It should be noted that we are always predicting

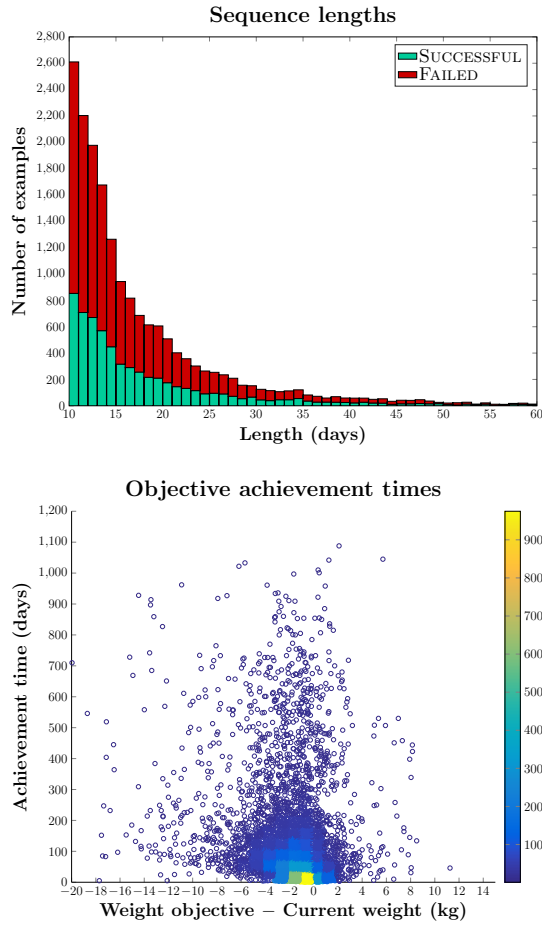


Figure 1: **Top:** Plot of the sequence length distribution in the final dataset. **Bottom:** Mixed heatmap/scatter plot of the weight objectives against their achievement times, for the successful sequences in the final dataset.

future trends from *historical information*—if a user successfully achieves their weight objective, the measurement that triggered this achievement will **not** be present in the given input sequence. Within our dataset, 6313 of the sequences represent successful examples, while the remaining 11723 represent examples of failure. To address the potential issues of class imbalance, appropriate class weights are applied to all optimisation targets and loss functions.

In order to get an impression of the statistics present within the dataset, we have visualised the sequence length distributions (outliers removed for visibility, but still used in our analysis), as well as scatter plots of successful weight objective magnitudes against their achievement times. These are provided by Figure 1.

We perform a task of *probabilistic classification* on the filtered dataset: predicting success for the weight objective,

evaluated using crossvalidation (this corresponds to a typical *binary classification problem*).

### 3 MODELS UNDER CONSIDERATION

#### Baseline models

In order to ascertain the suitability of deep recurrent models on this task, we have compared them on the objective classification task against several common baseline approaches to time-series classification, as outlined in [29]. For this purpose, we have considered four such models: *Support Vector Machines* (SVMs) using the RBF kernel, *Random Forests* (RFs), *Gaussian Hidden Markov Models* (GHMMs) and (feedforward) *Deep Neural Networks* (DNNs). The hyperparameters associated with the baseline models have been optimised with a thorough hyperparameter sweep—on a separate validation set—as detailed below.

- For the SVM, we have performed a grid search on its two hyperparameters ( $C$  and  $\gamma$ ) in the range  $\gamma \in 2^{[-15,5]}$ ,  $C \in 2^{[-5,15]}$ , finding the values of  $\gamma = 2^{-13}$  and  $C = 2^9$  to work best.
- For the RF, we have performed a search on the number of trees to use in the range  $N \in [10, 100]$ , finding  $N = 50$  to work best.
- For the GHMM, we have performed a search on the number of nodes to use in the range  $N \in [3, 40]$ , finding  $N = 7$  to work best.
- For the DNN, we have optimised the number of hidden layers (keeping the number of parameters comparable to the recurrent models) in the range  $\ell \in [1, 10]$ , finding  $\ell = 5$  to work best. This implied that each hidden layer had  $N = 120$  neurons. All hidden layers apply the *rectified linear* (ReLU) activation [17], and are regularised using batch normalisation [8] and dropout [24] with  $p = 0.5$ . All other relevant hyperparameters (such as the SGD optimiser and batch size) are the same as for the recurrent models.

For all the non-sequential models (SVM, RF, DNN), we have performed a search on the number of most recent time steps to use in the range  $l \in [5, 100]$ , finding  $l = 10$  to perform the best. The SVM model has been augmented to produce probabilistic predictions (and thus enable its ROC-AUC metric to be computed) by leveraging Platt scaling [20].

Lastly, we have attempted to use a classical time-series forecasting technique (ARIMA) to predict whether the weight objective threshold will be met in the future—with the performance being substantially worse than any machine learning model utilised.

#### Long short-term memory (LSTM)

All of our recurrent models are based on the long short-term memory (LSTM) [7] model. A single LSTM layer consumes a

sequence,  $\vec{x}$ , and produces an output sequence,  $\vec{y}$ . It also has storage capabilities through a *memory cell*,  $c$ . The size of the memory cell determines the *number of features* computed by the LSTM for each step of the output sequence. The output is produced in a step-by-step manner, such that the current output component,  $y_t$ , is produced taking into account the current input component ( $x_t$ ), the previous output ( $y_{t-1}$ ) and the previous cell state ( $c_{t-1}$ ). A single LSTM processing step also updates the memory cell to a new value,  $c_t$ . The equations describing this step within the LSTM model we employed (similar to [5]) are as follows:

$$i_t = \tanh(\mathbf{W}^{xi}x_t + \mathbf{W}^{yi}y_{t-1} + b^i) \quad (1)$$

$$j_t = \tilde{\sigma}(\mathbf{W}^{xj}x_t + \mathbf{W}^{yj}y_{t-1} + b^j) \quad (2)$$

$$f_t = \tilde{\sigma}(\mathbf{W}^{xf}x_t + \mathbf{W}^{yf}y_{t-1} + b^f) \quad (3)$$

$$o_t = \tilde{\sigma}(\mathbf{W}^{xo}x_t + \mathbf{W}^{yo}y_{t-1} + b^o) \quad (4)$$

$$c_t = c_{t-1} \otimes f_t + i_t \otimes j_t \quad (5)$$

$$y_t = \tanh(c_t) \otimes o_t \quad (6)$$

In these equations,  $\mathbf{W}^*$  and  $b^*$  correspond to learnable parameters (weights and biases, respectively) of the LSTM layer, and  $\otimes$  corresponds to element-wise vector multiplication.  $\tanh$  is the hyperbolic tangent function, and  $\tilde{\sigma}$  is the *hard sigmoid function*:

$$\tilde{\sigma}(x) = \max(0, \min(1, 0.2x + 0.5)) \quad (7)$$

providing a fast piecewise-linear approximation to the standard logistic sigmoid function:

$$\sigma(x) = \frac{1}{1 + \exp(-x)} \quad (8)$$

To aid clarity, for the remainder of the model description, we will compress Equations 1–6 into

$$\text{LSTM}(\vec{x}) = \vec{y} \quad (9)$$

representing a single LSTM layer, with its internal parameters and memory cell state kept implicit.

Our primary architecture represents a three-layer deep LSTM model for processing the historical weight/sleep/steps data. After performing the LSTM operations, the features of the final computed LSTM output step are concatenated with the height, gender, age category and weight objective, providing the following feature representation:

$$\text{LSTM}(\text{LSTM}(\text{LSTM}(\vec{wt}||\vec{sl}||\vec{st})))_T||\text{ht}||\text{gdr}||\text{age}||\text{obj} \quad (10)$$

where  $\vec{wt}$ ,  $\vec{sl}$  and  $\vec{st}$  are the input features (for weight, sleep and steps, respectively),  $||$  corresponds to featurewise concatenation, and  $T$  is the length of the initial sequence. These features are passed through two fully connected neural network layers, connected to a single output neuron which utilises a logistic sigmoid activation.

The fully connected layers of the networks apply *rectified linear* (ReLU) activations. We initialise the LSTM weights

using Xavier initialisation [4], and its forget gate biases with ones [9]. Finally, the fully connected weights are initialised using He initialisation [6], as recommended for ReLUs. The models are trained for 200 epochs using the Adam SGD optimiser, with hyperparameters as described in [11], and a batch size of 1024. For regularisation purposes, we have applied batch normalisation to the output of every hidden layer and dropout with  $p = 0.1$  to the input-to-hidden transitions within the LSTMs [31].

### Cross-modal LSTM

For this task we also propose a novel *cross-modal LSTM* (*X-LSTM*) architecture which exploits the *multimodality* of the input data more explicitly in order to efficiently redistribute the LSTM's parameters. We initially partition the input sequence into three parts (sleep data, weight data, steps data), and pass *each of those* through a separate three-layer LSTM stream. We also allow for *information flow* between the streams in the second layer, by way of *cross-connections*, where features from a single sequence stream are passed and concatenated with features from another sequence stream (after being passed through an additional LSTM layer). Represented via equations, the computed outputs across the three streams are:

$$\vec{h}_1^{\{wt, sl, st\}} = \text{LSTM}(\{\vec{wt}, \vec{sl}, \vec{st}\}) \quad (11)$$

$$\vec{h}_2^{\{wt \rightarrow wt, sl \rightarrow sl, st \rightarrow st\}} = \text{LSTM}(\{\vec{h}_1^{wt}, \vec{h}_1^{sl}, \vec{h}_1^{st}\}) \quad (12)$$

$$\vec{h}_2^{\{wt \leftrightarrow sl, wt \leftrightarrow st\}} = \text{LSTM}(\{\vec{h}_1^{wt}, \vec{h}_1^{wt}\}) \quad (13)$$

$$\vec{h}_2^{\{sl \leftrightarrow wt, sl \leftrightarrow st\}} = \text{LSTM}(\{\vec{h}_1^{sl}, \vec{h}_1^{sl}\}) \quad (14)$$

$$\vec{h}_2^{\{st \leftrightarrow wt, st \leftrightarrow sl\}} = \text{LSTM}(\{\vec{h}_1^{st}, \vec{h}_1^{st}\}) \quad (15)$$

$$\vec{h}_3^{wt} = \text{LSTM}(\vec{h}_2^{wt \rightarrow wt} || \vec{h}_2^{sl \leftrightarrow wt} || \vec{h}_2^{st \leftrightarrow wt}) \quad (16)$$

$$\vec{h}_3^{sl} = \text{LSTM}(\vec{h}_2^{sl \rightarrow sl} || \vec{h}_2^{wt \leftrightarrow sl} || \vec{h}_2^{st \leftrightarrow sl}) \quad (17)$$

$$\vec{h}_3^{st} = \text{LSTM}(\vec{h}_2^{st \rightarrow st} || \vec{h}_2^{wt \leftrightarrow st} || \vec{h}_2^{sl \leftrightarrow st}) \quad (18)$$

We used  $\vec{h}_2^{\{x, y, z\}} = \text{LSTM}(\{a, b, c\})$  to denote  $\vec{h}_2^x = \text{LSTM}(a)$ ,  $\vec{h}_2^y = \text{LSTM}(b)$ ,  $\vec{h}_2^z = \text{LSTM}(c)$ .

Finally, the feature representation passed to the fully connected layers is obtained by concatenating the final LSTM frames across all of the three streams:

$$(\vec{h}_3^{wt} || \vec{h}_3^{sl} || \vec{h}_3^{st})_T || \text{ht} || \text{gdr} || \text{age} || \text{obj} \quad (19)$$

The illustration of the entire construction process from individual building blocks is shown in Figure 2. This construction is biologically inspired by *cross-modal systems* [3] within the visual and auditory systems of the human brain—wherein several cross-connections between various sensory networks have been discovered [1, 30]. Similar techniques have already been successfully applied for handling sparsity

within convolutional neural networks [27] and audiovisual data integration [2].

To provide breadth, we evaluate *three* cross-connecting strategies: one as described by Equations 11–18 (A), one where the cross-connection does not have intra-layer LSTMs (B), and one where we don't cross-connect at all (N). The latter corresponds the most to prior work on multimodal deep learning [10, 19, 25]. Note that the variant (N) allows for computing the largest number of features within the parameter budget out of all three variants—no parameters being spent on cross-connections. The three scenarios are illustrated by Figure 3.

Finally, a recent state-of-the-art approach in processing multimodal sequential data [21] imposes cross-modality by weight sharing between the different modalities' recurrent weights ( $W^{y^*}$  in Equations 1–4)—we will refer to this technique as SH-LSTM. This comes at a cost to expressivity—in order to share them, these weight matrices need to have the same sizes, implying the different modality streams need to all compute the *same number of features* at each depth level. Keeping the parameter count comparable to the baseline LSTM, we evaluate three strategies for weight sharing (Figure 3): sharing across all modalities (ALL) and sharing across weight/sleep only, with (WSL) and without (CUT) steps data. This has been motivated by the fact that the weight and sleep data have, on their own, been found to be significantly more influential than steps data—as will be discussed in the Results section.

### X-LSTM hyperparameter tuning

In practice, we anticipate X-LSTMs to be derived from a baseline LSTM, in order to redistribute its parameters more efficiently. However, X-LSTMs might introduce an overwhelming amount of hyperparameters to tune. To make the process less taxing, we focus on the *meaning* of the feature counts—their comparative values are supposed to track the *relative significance* of each modality. First, we attempt to solve the task with a basic LSTM architecture using *only one of the modalities*. When scores (e.g. accuracies or AUC)  $s_{wt}$ ,  $s_{sl}$  and  $s_{st}$  are obtained for all three modalities, we redistribute the intra-layer feature counts of the X-LSTM according to the ratio  $s_{wt} : s_{sl} : s_{st}$ . This presents a good starting point for setting up constraints on the cross-connection ( $x \rightsquigarrow y$ ) feature counts, where the relative performance of the modalities  $x$  and  $y$  is reflected on the number of features sent across them.

However, this strategy will tend to create too “uniform” networks if the unimodal performance metrics are too similar, as it was the case with our experiments. To enforce larger discrepancies, we raise the obtained scores to a power  $k$ . This controls the tendency of the network to favour the most predictive modality when redistributing features. For a




Figure 2: A hierarchical illustration of a deep X-LSTM model with three layers and one cross-connection in the second layer. **Top:** A single LSTM block; all intermediate results, as described in Equations 1–6 ( $i_t$ ,  $j_t$ ,  $f_t$  and  $o_t$ ) are clearly marked. **Middle:** Replicating the LSTM cell to create an LSTM layer (for processing a given input sequence  $\vec{x}$ ). **Bottom:** A cross-modal deep LSTM model with two streams of three layers, taking sequences of length 3. In the second layer, the hidden sequences are shared between the two streams by being passed through a separate LSTM layer and feature-wise concatenated with the main stream hidden sequence.

fixed choice of  $k$ , we solve a system of equations in order to derive feature counts for all the intra-layer LSTM layers in an X-LSTM. Thus, most of the effort amounts to finding just *one* hyperparameter— $k$ . This allowed us to systematically obtain rapid improvements on the baseline, significantly shortening the period of trial-and-error with exploring the full hyperparameter space.




Figure 3: The three types of cross-connection strategies, and the three types of weight sharing strategies. Each arrow is an LSTM layer, dashed lines indicate the identity transformation, and all arrows going into the same node are concatenated. Connections between lines in the bottom row represent recurrent weight sharing. Top, left-to-right: X-LSTM (A), X-LSTM (B), X-LSTM (N). Bottom, left-to-right: SH-LSTM (ALL), SH-LSTM (WSL), SH-LSTM (CUT).

## 4 RESULTS

### Weight objective success classification

We performed stratified 10-fold crossvalidation on the baseline classifiers as well as the proposed LSTM model. Given the bias of the obtained data towards failure (there being *twice as many* sequences labelled unsuccessful), and the fact that it is not generally obvious what the classification threshold for this task should be (it likely involves several tradeoffs), we use **ROC curves** (and the associated *area* under them) as our primary evaluation metric. For completeness, we also report the accuracy, precision, recall,  $F_1$  score and the Matthews Correlation Coefficient [16] under the classification threshold which maximises the  $F_1$  score.

Afterwards we sought to construct competitive X-LSTMs, and therefore we computed the AUCs of the individual unimodal LSTMs on a validation dataset, obtaining AUCs of 80.62% (for weight), 80.17% (for sleep) and 74.18% (for steps). As anticipated, this was not far enough in order to reliably generate non-uniform X-LSTMs, so we proceeded to perform a grid search on the parameter  $k$ . We’ve originally taken steps of 5, but as we found the differences between adjacent steps to be negligible, we report the AUC results for  $k \in \{10, 20, 30\}$ . The X-LSTM performed the best with  $k = 30$ , and (B) cross-connections—we compare it directly with the LSTM, as well as the SH-LSTMs, and report its architecture in Table 1.

LSTM	X-LSTM (B, $k = 30$ )
76377 param.	75089 param.
21 features	wt: 15 features, sl: 12 features, st: 2 features <b>wt <math>\rightsquigarrow</math> sl: 9 features, wt <math>\rightsquigarrow</math> st: 14 features</b> <b>sl <math>\rightsquigarrow</math> wt: 6 features, sl <math>\rightsquigarrow</math> st: 11 features</b> <b>st <math>\rightsquigarrow</math> wt: 1 feature, st <math>\rightsquigarrow</math> sl: 1 feature</b>
42 features	wt: 29 features, sl: 24 features, st: 3 features
84 features	wt: 57 features, sl: 48 features, st: 5 features
	Fully connected, 128-D
	Fully connected, 64-D
	Fully connected, 1-D

Table 1: Architectures for the considered LSTM and cross-modal LSTM models. Cross-connections are highlighted.

To confirm that the advantages demonstrated by our methodology are statistically significant, we have performed paired  $t$ -testing on the metrics of individual cross-validation folds, choosing a significance threshold of  $p < 0.05$ . We find that all of the observed advantages in ROC-AUC are indeed statistically significant—verifying simultaneously that the recurrent models are superior to other baseline approaches, that the X-LSTM has significantly improved on its LSTM baselines and that cross-connecting is statistically beneficial (given the weaker performance of X-LSTM (N) despite being able to compute the most features overall). The SH-LSTM performed the best in its (WSL) variant (which allowed for more features to be allocated to weight and sleep streams, at the expense of the steps stream) but was even then unable to outperform the baseline LSTM—highlighting once again its lack of ability to accurately specify relative importances between modalities, which is essential for this task. The results are summarised by Tables 2–3 and Figure 4.


### Weight objective magnitude effects

An obvious factor that will affect the model’s performance is the *magnitude* of weight objectives that the users are setting. To illustrate the X-LSTM’s predictive power, we have aggregated its predictions across all of the crossvalidation folds (for a classification threshold of 0.5). From these predictions, we have produced a histogram of the proportions of correctly classified, incorrectly classified successful and incorrectly classified failed sequences, for various bins of weight objective magnitudes (ref. Figure 5).

It should be expected that, at smaller weight objective magnitudes, the model’s confusions will be roughly *evenly distributed* between successes and failures. As objectives

Metric	ARIMA	SVM	RF	GHMM	DNN	LSTM	SH-LSTM	X-LSTM
Accuracy	59.22%	67.65%	70.97%	66.31%	68.93%	79.12%	78.49%	<b>80.30%</b>
Precision	42.96%	52.54%	56.05%	51.26%	53.80%	67.25%	65.31%	<b>68.66%</b>
Recall	50.44%	81.02%	81.34%	82.32%	<b>83.02%</b>	79.30%	82.95%	81.62%
F <sub>1</sub> score	46.40%	63.71%	66.25%	63.11%	65.18%	72.69%	72.98%	<b>74.37%</b>
MCC	13.94%	39.74%	44.75%	38.57%	42.63%	56.60%	56.80%	<b>59.45%</b>
ROC AUC	—	76.77%	79.97%	74.86%	78.54%	86.91%	86.63%	<b>88.07%</b>
<i>p</i> -value	—	$2 \cdot 10^{-12}$	$6 \cdot 10^{-10}$	$7 \cdot 10^{-11}$	$2 \cdot 10^{-11}$	$1 \cdot 10^{-4}$	$4 \cdot 10^{-5}$	—

**Table 2: Comparative evaluation results of the baseline models against the LSTMs after 10-fold crossvalidation. Reported X-LSTM is the best-performing (B,  $k = 30$ ). Reported SH-LSTM is the best-performing (WSL). All metrics except the ROC AUC reported for the classification threshold that maximises the F<sub>1</sub> score. Reported *p*-values are for the X-LSTM vs. each baseline for the ROC-AUC metric.**



**Figure 4: Mean ROC curves for the baselines, LSTM and the best-performing SH-LSTM and X-LSTM models.**

become more ambitious, the model should become increasingly skewed towards *predicting failure more often* (as a consequence of the realistic statistics on such objectives), and therefore the majority of confusions are expected to represent successful sequences that the model misclassified as failures.

The results completely match these expectations—at lower weight objective magnitudes, there is a roughly 50/50 split between misclassified successes and failures. However, starting at  $-3\text{kg}$  and moving towards more ambitious objectives, there is a clear bias towards misclassifying successful sequences, which eventually grows into nearly all misclassified sequences being successful. This kind of behaviour is likely to be desirable—on average, it will encourage the users to

Model	$k = 10$	$k = 20$	$k = 30$
X-LSTM (A)	<b>87.60%</b>	<b>87.60%</b>	87.75%
X-LSTM (B)	87.21%	87.56%	<b>88.07%</b>
X-LSTM (N)	86.49%	86.98%	87.30%
<i>p</i> -value	$9.55 \cdot 10^{-5}$	0.021	$1.03 \cdot 10^{-3}$
SH-LSTM (ALL)	85.58%		
SH-LSTM (WSL)	86.63%		
SH-LSTM (CUT)	86.30%		

**Table 3: Effects of varying the hyperparameter  $k$  and cross-connecting strategy of X-LSTMs to the mean ROC AUC after crossvalidation. Reported *p*-values are for the (N) vs.  $\max(A, B)$  strategies. We also report the mean ROC AUC for the three kinds of sharing strategies of SH-LSTMs.**

choose realistic weight objectives, at the expense of making incorrect initial predictions about a few users that do eventually manage to achieve very ambitious goals.

### Visualising detected features

It is hard to interpret the parameters of a network directly, so instead we focus on *generating artificial sequences that maximise the network’s confidence in success or failure* [23]. Specifically, we apply the approach of *visualising classification models*, which is now a standard technique in the computer vision community [23]. Starting from a zero-input sequence,  $\mathbf{I}_0 = \mathbf{0}$ , we would like to produce an input  $\mathbf{I}'$  that maximises the classification confidence:

$$\mathbf{I}' = \underset{\mathbf{I}}{\operatorname{argmax}} \Sigma(\mathbf{I}) - \lambda \|\mathbf{I}\|_2^2 \quad (20)$$

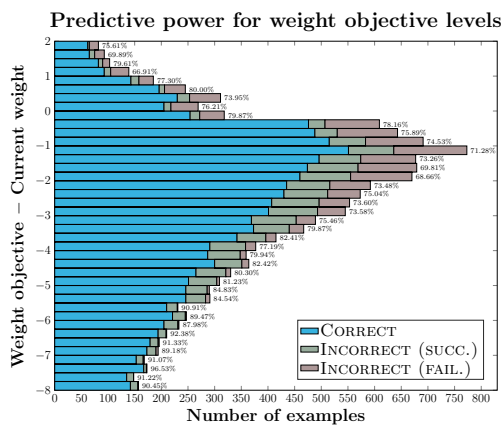



Figure 5: Top: A bar plot demonstrating the X-LSTM’s performance for different magnitudes of weight objectives (at the classification threshold of 0.5). Bottom: The same plot, zoomed in on the  $[-8, 2]$  range of weight objectives (where the majority of the examples are).

where  $\Sigma(I)$  is the network’s output for  $I$ , and  $\lambda$  is an  $L_2$ -regularisation parameter (to penalise large day-to-day vari-ances). We found that  $\lambda = 5$  works best.

Generated sequences spanning 10 days are shown in Fig. 6. As expected, we observe that a user is likely to hit their weight objective if there is a downwards trend in weight and an upwards trend in steps, and vice-versa for a failing sequence. Interestingly, the model also uncovered that to have a higher confidence of success, it is important for the user to *fall asleep quicker once going to bed*. This is likely encoding important *latent variables* that we can not directly access from the dataset—for example, a person that takes more time to fall asleep is more likely to snack in the evening, which is known to be detrimental to weight loss (as previously observed in biomedical research [12, 18, 22]).

Aside from helping us establish a more illustrated notion of the neural network’s decision-making process, this analysis may also enhance *feature selection*. Namely, we may consider


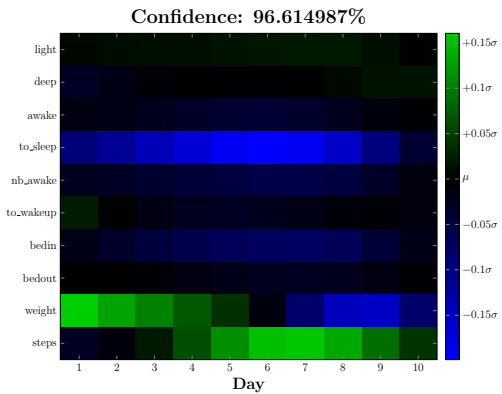


Figure 6: Iteratively produced artificial sequences that maximise the model’s confidence in achieving (top) or failing (bottom) a  $-4\text{kg}$  weight objective. Best viewed in colour.

the features that are not significantly differently activated in the two cases to be *irrelevant*—here, those are likely to be the time spent awake during the night, number of times the person is awoken during the night, and the bed-in/bed-out times.

## 5 CONCLUSION

In this work, we applied RNNs to fitness data with the aim of inferring the probability of users achieving their weight objective. Our novel *cross-modal LSTM* (X-LSTM) achieves the best performance by exploiting the multimodality in this data, resulting in higher ROC-AUC performance on this task than any other model considered—including a previous state-of-the-art approach to cross-modal sequential data processing—with statistical significance. Our results show the viability of a new concrete application of learning within consumer health care, despite the inherent noise and sparsity present in such data. More broadly, our X-LSTM architecture hints at a new approach to modelling multimodal time-series data in general.

## REFERENCES

- [1] Anton L. Beer, Tina Plank, and Mark W. Greenlee. 2011. Diffusion tensor imaging shows white matter tracts between human auditory and visual cortex. *Experimental Brain Research* 213, 2 (2011), 299–308. <https://doi.org/10.1007/s00221-011-2715-y>
- [2] Cătălina Cangea, Petar Veličković, and Pietro Liò. 2017. XFlow: 1D-2D Cross-modal Deep Neural Networks for Audiovisual Classification. *arXiv preprint arXiv:1709.00572* (2017).
- [3] Mark A Eckert, Nirav V Kamdar, Catherine E Chang, Christian F Beckmann, Michael D Greicius, and Vinod Menon. 2008. A cross-modal system linking primary auditory and visual cortices: Evidence from intrinsic fMRI connectivity analysis. *Human brain mapping* 29, 7 (2008), 848–857.
- [4] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*. 249–256.
- [5] Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850* (2013).
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*. 1026–1034.
- [7] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [8] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [9] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. 2015. An Empirical Exploration of Recurrent Network Architectures. In *Proceedings of The 32nd International Conference on Machine Learning*. 2342–2350.
- [10] Lukasz Kaiser, Aidan N Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017. One model to learn them all. *arXiv preprint arXiv:1706.05137* (2017).
- [11] Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [12] C Kleiser, N Wawro, M Stelmach-Mardas, H Boeing, K Gedrich, H Himmerich, and J Linseisen. 2017. Are sleep duration, midpoint of sleep and sleep quality associated with dietary intake among Bavarian adults? *European Journal of Clinical Nutrition* (2017).
- [13] Neal Lathia, Veljko Pejovic, Kiran K Rachuri, Cecilia Mascolo, Mirco Musolesi, and Peter J Rentfrow. 2013. Smartphones for Large-Scale Behavior Change Interventions. *IEEE Pervasive Computing* 12, 3 (2013), 66–73.
- [14] Anming Li, Hareton Leung, and Yvette Lui. 2014. Friend Recommendation for Weight Loss App. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct)*. ACM, New York, NY, USA, 1257–1264. <https://doi.org/10.1145/2638728.2638807>
- [15] Edith Talina Luhanga, Akpa Akpro Elder Hippocrate, Hirohiko Suwa, Yutaka Arakawa, and Keiichi Yasumoto. 2016. Towards Proactive Food Diaries: A Participatory Design Study. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 1263–1266. <https://doi.org/10.1145/2968219.2968441>
- [16] Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure* 405, 2 (1975), 442–451.
- [17] Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. 807–814.
- [18] Arlet V Nedeltcheva, Jennifer M Kilkus, Jacqueline Imperial, Kristen Kasza, Dale A Schoeller, and Plamen D Penev. 2009. Sleep curtailment is accompanied by increased intake of calories from snacks. *The American journal of clinical nutrition* 89, 1 (2009), 126–133.
- [19] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*. 689–696.
- [20] John C. Platt. 1999. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. In *ADVANCES IN LARGE MARGIN CLASSIFIERS*. MIT Press, 61–74.
- [21] Jimmy Ren, Yongtao Hu, Yu-Wing Tai, Chuan Wang, Li Xu, Wenxiu Sun, and Qiong Yan. 2016. Look, Listen and Learn – a Multimodal LSTM for Speaker Identification. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, 3581–3587. <http://dl.acm.org/citation.cfm?id=3016387.3016407>
- [22] Natsuko Sato-Mito, Satoshi Sasaki, Kentaro Murakami, Hitomi Okubo, Yoshiko Takahashi, Shigenobu Shibata, Kazuhiko Yamada, Kazuto Sato, Freshmen in Dietetic Courses Study II Group, et al. 2011. The midpoint of sleep is associated with dietary intake and dietary behavior among young Japanese women. *Sleep medicine* 12, 3 (2011), 289–294.
- [23] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013).
- [24] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [25] Nitish Srivastava and Ruslan R Salakhutdinov. 2012. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*. 2222–2230.
- [26] Petar Veličković, Laurynas Karazija, Nicholas D Lane, Sourav Bhattacharya, Edgar Liberis, Pietro Liò, Angela Chieh, Otmane Bellahsen, and Matthieu Vegreville. 2017. Cross-modal Recurrent Models for Human Weight Objective Prediction from Multimodal Time-series Data. *arXiv preprint arXiv:1709.08073* (2017).
- [27] Petar Veličković, Duo Wang, Nicholas D Lane, and Pietro Liò. 2016. X-CNN: Cross-modal Convolutional Neural Networks for Sparse Datasets. *arXiv preprint arXiv:1610.00163* (2016).
- [28] Sinead Watson, V. Jayne Woodside, J. Lisa Ware, J. Steven Hunter, Alanna McGrath, R. Christopher Cardwell, M. Katherine Appleton, S. Ian Young, and C. Michelle McKinley. 2015. Effect of a Web-Based Behavior Change Program on Weight Loss and Cardiovascular Risk Factors in Overweight and Obese Adults at High Risk of Developing Cardiovascular Disease: Randomized Controlled Trial. *J Med Internet Res* 17, 7 (16 Jul 2015), e177. <https://doi.org/10.2196/jmir.3828>
- [29] Zhengzheng Xing, Jian Pei, and Eamonn Keogh. 2010. A brief survey on sequence classification. *ACM SIGKDD Explorations Newsletter* 12, 1 (2010), 40–48.
- [30] Weiping Yang, Jingjing Yang, Yulin Gao, Xiaoyu Tang, Yanna Ren, Satoshi Takahashi, and Jinglong Wu. 2015. Effects of Sound Frequency on Audiovisual Integration: An Event-Related Potential Study. *PLoS ONE* 10, 9 (09 2015), 1–15. <https://doi.org/10.1371/journal.pone.0138296>
- [31] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329* (2014).