Robotica

http://journals.cambridge.org/ROB

Additional services for Robotica:

Email alerts: <u>Click here</u> Subscriptions: <u>Click here</u> Commercial reprints: <u>Click here</u> Terms of use : <u>Click here</u>



A novel approach for representing and generalising periodic gaits

Hsiu-Chin Lin, Matthew Howard and Sethu Vijayakumar

Robotica / Volume 32 / Special Issue 08 / December 2014, pp 1225 - 1244 DOI: 10.1017/S026357471400188X, Published online: 13 August 2014

Link to this article: http://journals.cambridge.org/abstract_S026357471400188X

How to cite this article:

Hsiu-Chin Lin, Matthew Howard and Sethu Vijayakumar (2014). A novel approach for representing and generalising periodic gaits. Robotica, 32, pp 1225-1244 doi:10.1017/S026357471400188X

Request Permissions : Click here



A novel approach for representing and generalising periodic gaits

Hsiu-Chin Lin^{†*}, Matthew Howard[‡] and Sethu Vijayakumar[†]

†Institute of Perception Action and Behaviour, School of Informatics, University of Edinburgh, Edinburgh, UK ‡Department of Informatics, King's College London, London, UK

(Accepted June 19, 2014. First published online: August 13, 2014)

SUMMARY

Our goal is to introduce a more appropriate method of representing, generalising and comparing gaits; particularly, walking gait. Human walking gaits are a result of complex, interdependent factors that include variations resulting from embodiments, environment and tasks, making techniques that use average template frameworks suboptimal for systematic analysis or corrective interventions. The proposed work aims to devise methodologies for being able to represent gaits and gait transitions such that optimal policies that eliminate the inter-personal variations from tasks and embodiments may be recovered. Our approach is built upon (i) work in the domain of nullspace policy recovery and (ii) previous work in generalisation for point-to-point movements. The problem is formalised using a walking-phase model, and the nullspace learning method is used to generalise a consistent policy from multiple observations with rich variations. Once recovered, the underlying policies (mapped to different gait phases) can serve as reference guideline to quantify and identify pathological gaits while being robust against interpersonal and task variations. To validate our methods, we have demonstrated robustness of our method with simulated sagittal two-link gait data with multiple ground truth constraints and policies. Pathological gait identification was then tested on real-world human gait data with induced gait abnormality, with the proposed method showing significant robustness to variations in speed and embodiment compared to template-based methods. Future work will extend this to kinetic features and higher dimensional features.

KEYWORDS: Rehabilitation; Human biomechanics; Robots for clinical assessment; Imitation; Bipeds.

1. Introduction

Many everyday human skills can be considered as a form of periodic movement. For example, locomotion can be considered as a periodic motion of the legs, and wiping a table can be a kind of periodic motion of the arms. The form of these movements is influenced by many factors, namely, (i) the embodiment of the subject (e.g., limb lengths, mass properties, etc.), (ii) the environment in which the behaviour is performed (e.g., is the subject walking on flat or uneven terrain?) and (iii) task contextual factors (e.g., is the subject hurrying to a meeting or just taking a walk in the park?). Nevertheless, despite these variations, some consistency appears that causes us to identify behaviours, such as walking, as belonging to the same class.

The fact that such variations exist within a single class of behaviour, such as walking, indicates the presence of redundancy in the system. That is, the presence of additional degrees of freedom allows the constraints induced by these various factors to be satisfied, while at the same time satisfying some underlying consistent behavioural goal. The latter could be to minimise effort, maintain comfort, or other such criteria. This dependence on the various factors makes modelling human gait hard in

* Corresponding author. E-mail: H.Lin-21@sms.ed.ac.uk





Fig. 1. We hypothesise that the behaviours that we observe are the combinations of some consistent characteristics and variations in embodiment and task factors. Our approach is based on examining various behaviours to see if such consistency can be found.

general,^{1,2} especially given the fact that the precise influence of different factors on the movement can be hard to assess (e.g., how is the foot placement strategy affected by differences in terrain^{3,4}).

Examples in modelling/comparing human gaits can be found in exoskeleton systems such as the Lokomat, the Skywalker and the Lower Extremity Powered Exoskeleton (LOPES).^{5–7} Although these devices are cleverly designed, much work is needed to personalise gait correction. One current issue is that existing systems normally restrict the patients to follow some predefined walking patterns, such that the patients are required to walk at a certain speed, slope or step-size. However, clinical results show that motivating the patients to walk more proactively at a preferred pace promotes the overall results of rehabilitation.⁸ Also, the pre-defined walking patterns are normally derived by taking an average gait template from healthy subjects, but it seems inadequate. For instance, this approach would consider faster or slower walks as deviations from a normal gait.

Some existing devices have a more flexible approach by incorporating impedance control and/or tolerating small deviation from the reference gait.^{9–11} However, taking the average template framework as the reference policy is still problematic. Instead of restricting the patients to follow these predefined rules of training, a more appropriate way is to let the patients walk the way they prefer and correct them only if needed. Our approach is different from the previous work in that we aim to extract the consistent aspect of normal gaits and separate those from the natural variations (from embodiments, environments and behaviours).

Our assumption is built upon previous research in gait analysis. Studies have shown very small variations in human kinematics when performing the same gait. For example, Stokes *et al.*¹² evaluated the repeatability of kinematic data of 40 subjects. Although there was variance in the range of motion, the kinematic patterns in the saggital plane were highly predictable. Ivanenko *et al.*¹³ applied dimensional reduction analysis to determine whether the pattern of muscle activities can be described by some underlying manifold. Their work showed that while the EMG (electromyography) signals varied significantly when walking at different speeds, using Principle Component Analysis,¹⁴ just the first five principle components of these EMGs could account for the main features of the signals.

In recent years, a number of new tools have become available in the learning and robotics community that allow data from constrained and/or redundant systems^{15,16} to be used to uncover underlying consistent behaviours that may be otherwise masked by the constraints. Our approach is based on examining behaviours in the light of such methods, to see if certain underlying characteristics of gait can be found in the face of variations (see Fig. 1). These may be interpersonal (e.g., one person walks with a particular style that maintains that person's most comfortable posture) or intrapersonal (e.g., it is likely that all people walk to minimise some measure of effort).

In this paper, we narrowed down our problem to a single class of human locomotion and analyse walking gaits on even terrain. We examine various walking behaviours within this class and see if any consistency can be found across persons. Potential applications of our work include rehabilitation, for example, through use of robotic support systems, and gait abnormality detection and diagnosis.^{5,6,17}

2. Background and Related Work

2.1. Phase-based decomposition of walking

A common approach in the gait analysis literature,¹⁸ and one that we will follow in this paper, is to decompose walking into a series of *cycles* where one walking cycle is defined as the time between two consecutive occurrences of an event.

As an example in Fig. 2, the instant at which one heel strikes the ground can be used to demark the beginning of a cycle, that continues until the same heel strikes the ground again. Within a cycle, the







Fig. 2. Classical definition of gait phase and the phase division in our experiment (extracted and modified $from^{18}$).



Fig. 3. Correspondence between leg and two-link system.

gait can be further divided into a series of walking phases based on a series of events. The dynamics of each walking phase is different from one another. In order to deal with these variations in different phases along with inherent variations in embodiment and task-related factors, a sufficiently flexible model is required that can be used to decompose the movement. In Section 3, we outline such a model based on task/nullspace decomposition of behaviour, and describe how it may be used to capture these variations in a unified way.

2.2. Kinematic representations of human gait

A simple model of the human leg can be described by a planar two-link system (Fig. 3). In this model, L_1 , L_2 are the lengths of femur and tibia, respectively, angles q_1 , q_2 represent the hip and knee angles, respectively, and r_1 , r_2 represent the horizontal and vertical positions of the foot.

For kinematic control, the state of the system can be represented by the joint angles $\mathbf{x} = (q_1, q_2)^{\top}$, and the controls can be the joint velocities $\mathbf{u} = (\dot{q}_1, \dot{q}_2)^{\top}$. In addition to working with purely kinematic representations, several previous publications have focused on using Zero Moment Point (ZMP) to model humanoid locomotion.^{19–21} Also popular are dynamical-system-based representations of periodic movements, since they have significant merits over time-indexed representations.²² The work of Morimoto and Atkeson²³ is based on a periodic pattern generator and a coupled oscillator model to modulate the phase of sinusoidal patterns, in addition to representations that are derived from central pattern generators (CPGs).^{24,25} In this paper, we will focus on the kinematics representations (see Fig. 3) with the aim to explore more degrees of freedom as well as kinetic features in subsequent investigations.

3. Decomposed Walking Model

In this section, we describe a constrained tracking control scheme in which behaviour is decomposed into task/nullspace components. We then describe our interpretation of walking gait and how to determine the control variables in this control scheme. To aid the explanation, we use the two-link system in Fig. 3 as an illustrative example.

3.1. Walking-phase models

We assume that movements within each walking phase are the result of a composition of components handling the *phase-critical* components of motion (i.e., those that must be controlled for successful completion of a given phase), and redundant components that control consistent aspects of motion. In order to satisfy the phase-critical components, it is assumed the set of constraints

$$\mathbf{A}_{k}(\mathbf{x},t)\,\mathbf{u}(\mathbf{x},t) = \mathbf{b}_{k}(\mathbf{x},t) \tag{1}$$

is maintained, where $\mathbf{x} \in \mathbb{R}^n$ represents state, $\mathbf{u} \in \mathbb{R}^m$ represents the action, *t* is time and *k* indexes the phase. Here, the *task-space policy* $\mathbf{b}_k(\mathbf{x}, t) \in \mathbb{R}^p$ (p < m) describes a *task-dependent* control policy. The *constraint matrix* $\mathbf{A}_k(\mathbf{x}, t) \in \mathbb{R}^{p \times m}$ is a matrix projecting the task-space policy onto the relevant part of the control space. Inverting Eq. (1) results in the relation

$$\mathbf{u}(\mathbf{x},t) = \mathbf{A}_k(\mathbf{x},t)^{\dagger} \, \mathbf{b}_k(\mathbf{x},t) + \mathbf{N}_k(\mathbf{x},t) \, \boldsymbol{\pi}(\mathbf{x}), \tag{2}$$

where \mathbf{A}^{\dagger} is the Moore–Penrose pseudo-inverse of \mathbf{A} , and we define $\mathbf{N}_k(\mathbf{x}, t) := (\mathbf{I} - \mathbf{A}_k(\mathbf{x}, t)^{\dagger} \mathbf{A}_k(\mathbf{x}, t)) \in \mathbb{R}^{m \times m}$ where $\mathbf{I} \in \mathbb{R}^{m \times m}$ is the identity matrix. In Eq. (2), the second term arises due to the redundancy (since p < m), and allows secondary control objectives to be realised through the control policy $\pi(\mathbf{x}) \in \mathbb{R}^m$.

We assume that \mathbf{A}_k and \mathbf{b}_k are not explicitly known, but the quantities vary across walking phases to handle different phase-critical components. We assume that the nullspace policy π is *independent* of the phase, however, the *observed effects* of control toward these objectives (i.e., the nullspace component of motion $\mathbf{N}_k \pi$) may be influenced by the phase. This is because the nullspace policy π is subject to the higher priority constraints imposed by \mathbf{A}_k and \mathbf{b}_k .

3.2. Decomposition of walking gait with a two-link system

We assume that the walking gait is a combination of *characteristics* of the gait and *variations* resulting from different embodiments and behaviours. In this section, we describe how to formulate the characteristics and variations with our walking-phase model.

3.2.1. Consistent characteristics of gaits. As previously stated, we hypothesise that there exists some consistent characteristics within a single class of human locomotion. In walking, such consistency may include energy minimisation, or maintenance of a comfort posture.

The consistent characteristics of walking is captured in our model as the underlying nullspace policy π . An example of such a policy can be a limit cycle policy $\dot{r} = r(\rho - r^2)$, where r and θ are the polar representations of the state such that $\mathbf{x} = (r \cos(\theta), r \sin(\theta)), \rho$ is the radius of the attractor and $\dot{\theta}$ is the angular velocity.

3.2.2. Variations in embodiments. The variations in embodiments arise from inter-personal differences in factors such as body size, body type and physical limits. In our model, we assume that such variations will result in modifications of the *constraint matrix* **A**. (Hence, **A** varies across phases and subjects.)

In Eq. (1), we define the constraint matrix $\mathbf{A} \in \mathbb{R}^{p \times m}$ as a set of *p* task constraints, and each task constraint refers to restrictions on the freedom of some subspace of the system. The task constraints can be imposed in different representations (i.e., joint-space, end-effector space or other sophisticated transformation).

A simple case is to restrict some sub-space of the joint-space. For example, if the state of the system is defined as $\mathbf{x} = (q_1, q_2)^{\top}$, and the constraint matrix is set as $\mathbf{A} = [0, 1]$, the knee angle is restricted to follow various task-space policy **b** while the hip angle is free to move with the nullspace policy $\boldsymbol{\pi}$.

Another example of a constraint is to restrict the end-effector position $\mathbf{r} = [r_1, r_2]$. For instance, by observing the hip position in the sagittal plane when walking at various speeds, the horizontal dimension has a monotonic motion forward at various velocities, while the vertical dimension consistently moves up and down, regardless of the speeds. One way of characterising such motion



Mode

Í

 π

Controlled variables		Affec	Affected variables	
Variation		Example	b	Α
behaviou	r	speed, cadence	\checkmark	
embodim	ents	leg-length		\checkmark
environm	ent	slopes		\checkmark
phase		-		\checkmark

Variations in

b

u

Table I. Correspondence between variations in walking gaits and the variables in our proposed model.

Fig. 4. (Colour online) Correspondence between walking gaits and walking-phase model.

Variations in

А

would be to define task constraints, the Jacobian and the constraint matrix as

$$\hat{\mathbf{n}} = \begin{bmatrix} 1 & 0 \end{bmatrix}, \mathbf{J} = \begin{bmatrix} \frac{\partial r_1}{\partial q_1} & \frac{\partial r_1}{\partial q_2} \\ \frac{\partial r_2}{\partial q_1} & \frac{\partial r_2}{\partial q_2} \end{bmatrix}, \quad \mathbf{A} = \hat{\mathbf{n}} \mathbf{J} = \begin{bmatrix} \frac{\partial r_1}{\partial q_1} & \frac{\partial r_1}{\partial q_2} \end{bmatrix}$$

 $N\pi$

This would allow motion in the vertical space to follow the consistent nullspace policy provided that the horizontal dimension moves as required by the tasks.

3.2.3. Variations in environments. Another variation of walking gait comes from factors such walking terrain and slope, collectively aggregated in our model as variation in the environment. In our model, they are also captured by variation of the constraint matrix **A**.

3.2.4. Variations in behaviours. The variations in behaviours are caused by other contextual factors such as the need to hurry for a meeting, and result in variations in, for example, step-sizes and speeds. In our model, these variations are captured by changes to the *task-space policy* **b**, and we assume that **b** may vary across gait cycles to handle different behaviours.

In Eq. (1), we define the task-space policy $\mathbf{b} \in \mathbb{R}^p$ outputs the task-space velocity, in order to accomplish some operations. Depending on which dimension is constrained (i.e., defined by **A**), the constrained dimensions are restricted to move to a specific target along that dimension.

For instance, in the terminal swing phase, the foot is placed on the ground at the end of the phase. **A** may be defined such that the foot position is controlled, while **b** may describe a point attractor driving the foot to a desired placement. An example of **b** could be a point-attractor in task space $\mathbf{b}(\rho) = \omega(\rho^* - \rho)$, where ω is a parameter that controls the speed, ρ^* is the task-space target, and ρ is the task-space state. The position of heel strike depends on the step-length of that cycle; in this case, ρ^* can vary to handle different step lengths.

Table I summarises how the variations in embodiments and behaviours correspond to the control parameters, and how these parameters affect the variables in the model.

Figure 4 shows the correspondence between walking gaits and the walking-phase model. The observed behaviours **u** are the result of some consistent policy π modulated by various **A**, **b**. With this formulation, we examine various **u** and see if π can be recovered.

Figure 5 illustrates examples of behaviours produced from three different task constraints. We set up a simple system consisting of a nullspace policy $\pi(t) = [\cos(t), \sin(t)]$ and a task-space policy $\mathbf{b}(t) = \alpha \times \sin(0.25t + \alpha)$, where α was drawn from $\alpha \sim U[.4, 1]$. If the constraint is set to A = [0, 0] (fully unconstrained, Fig. 5(a)), the movements are simply the output of the nullspace policy π . If the constraint is set to $\mathbf{A} = [0, 1]$ (Fig. 5(b)), the hip angle is unconstrained, and the knee





Fig. 5. (Colour online) Examples of three different task constraints.

angle is constrained to follow the task-space policy **b**, resulting in variations driven by the task. If the constraint is set to $\mathbf{A} \approx [0.7, 0.7]$ (Fig. 5(c)), the observations are the combination of the nullspace policy π and the task-space policy **b**, and both hip and knee show variance in behaviours consistent with the task.

3.3. An example of modelling walking phases

Here, we describe an example of modelling the phase-critical components of three walking phases in Fig. 2. We utilise the interpretation and results from ref. [26] to choose **A** and **b** for each phase.

Note that using three phases is not conventional and not necessarily optimal. Our method works when A_k are consistent within a phase, and the least number of subsets is more preferable. We are interested in analysing these three phases since they seem to be more different from one another.

3.3.1. *Phase 1: stance phase.* The hip flexes to rotate backward and then reverse the rotation right before the end of stance phase. The leg is almost straight so that the stance leg is able to support the body weight.

This condition seems required for all kinds of walking behaviours. We consider that there is some level of redundancy/consistency in both hip and knee angles, so we set the constraint as $\mathbf{A} = [-0.6, 0.8]$. The task-space target is set to $\rho^* = (-130^\circ, -10^\circ)$ since this is a reasonable approximation of the posture at the end of stance phase.

This phase ends when the horizontal displacement between the heel and the torso is approximately 80% of the step-size. We utilise the initial horizontal position of the feet \mathbf{r}_1^0 and terminate the phase when $\mathbf{r}_1 < -0.8\mathbf{r}_1^0$.

3.3.2. *Phase 2: pre-swing to mid-swing point*. Our second phase combines pre-swing and initialswing phases. During these two phases, the foot lifts off from the ground as the hip joint extends and the knee joint flexes.

How fast the knee angle rises is dependent on the speed of walking, so we define the constraint as $\mathbf{A} = [-0.4, 0.9]$. At the end of phase 2, the hip joint is almost at its maximum extension and the knee joint is at its maximum flexion, so we set task-space target as $\rho^* = (-60\degree, -120\degree)$.

The knee angle should reach its maximum flexion at the end of this phase, so the phase terminates when $\mathbf{x}_2 < \mathbf{x}_2^{\min}$.

3.3.3. *Phase 3: mid-swing point to end of cycle.* The thigh is rotating forward and the knee is once again straightened.

The movement of the knee is highly dependent on both speed and step-size. (In addition to the speeds, where the heel strikes on the ground would also depend on the step-sizes.) Therefore, the constraint is set to $\mathbf{A} = [0.1, 0.99]$. We assume that maximum hip extension is approximately -40° , and the maximum knee extension is 0° , so the task-space target is set to $\rho^* = (-40^{\circ}, 0^{\circ})$.

At the end of this phase, the foot is set back on the ground. We define the height of the left feet as \mathbf{r}_2 and the initial height of the left feet as \mathbf{r}_2^0 , and the terminal condition is set to $\mathbf{r}_2 < \mathbf{r}_2^0$.



Thank p Tor each phase.				
Phase	Constraint A	Task-space target ρ^*	End condition	
1	[-0.6, 0.8]	(-130°, -10°)	$\mathbf{r}_1 < -0.8\mathbf{r}_1^0$	
2	[-0.4, 0.9]	(-60°, -120°)	$\mathbf{x}_2 < \mathbf{x}_2^{\min}$	
3	[0.1, 0.99]	(-40°, 0°)	$\mathbf{r}_2 < \mathbf{r}_2^0$	

Table II. An example of modelling phase-critical components by defining parameters A and ρ^* for each phase.

Under our constraint model, the constraint matrix **A**, the task-space target ρ^* and the terminal condition for each phase are summarised in Table II.

4. Generalising Periodic Gaits

In this section, we describe our approach to generalising the characteristics of walking gaits by reconstructing the unconstrained policy.

Our method works on data given as N pairs of observed states **x** and observed actions **u**. We assume that (i) the observations can be decomposed as $\mathbf{u} = \mathbf{A}^{\dagger} \mathbf{b} + \mathbf{N} \boldsymbol{\pi}$, (ii) **u** are generated using the same nullspace policy $\boldsymbol{\pi}$, (iii) each observation might have been constrained to accomplish some tasks (that is, $\mathbf{A} \mathbf{u} = \mathbf{b}$ for some constraint matrix $\mathbf{A} \neq 0$ and task-space policy $\mathbf{b} \neq 0$) and (iv) **b** and **A** (and **N**) are not explicitly known for any given observation. The goal of learning is to approximate the policy $\boldsymbol{\pi}$, that characterises the gait, and is independent of task and embodiment.

4.1. Naive (average template) approach

Given the observations \mathbf{x} , \mathbf{u} , a simple but naive approach to modelling the data is to apply direct regression to estimate the policy function. Note that this corresponds to the default *average template* solution used in many gait and rehabilitation analyses. Specifically, one may minimise the error

$$E_{\text{DPL}} = \sum_{n=1}^{N} \|\mathbf{u}_n - \tilde{\boldsymbol{\pi}}_n\|^2, \qquad (3)$$

where $\tilde{\pi}$ is some suitable estimator of the policy function. The estimator could be a linear estimator (such as linear regression) or a non-linear estimator (such as a network of radial basis functions (RBFs)²⁷).

Such an approach ignores the variations in constraints A and task-space policy b, and yields the average motion from different observations. However, this is unrealistic in everyday behaviour, so minimising Eq. (3) is unlikely to result in a good model of the gait.

4.2. Nullspace policy learning for periodic gaits

Considering the analysis in Section 3, in this paper we take an alternative approach, in which the constraint and task-space variations are explicitly considered. The proposed approach builds on previous research on policy recovery¹⁶ for point-to-point movements, and adapts it to generalise the characteristics of gaits. The key additional challenges in walking tasks are (i) the differences across subjects due to interpersonal variations and (ii) the temporal switching of constraints between phases.

In the proposed approach, the policy π is estimated using two separate steps. The first step is to decompose the observations **u** into two orthogonal components: the task-space component $\mathbf{u}^{ts} \equiv \mathbf{A}^{\dagger} \mathbf{b}$ and the nullspace component $\mathbf{u}^{ns} \equiv \mathbf{N} \pi$ such that $\mathbf{u} = \mathbf{u}^{ts} + \mathbf{u}^{ns}$. The second step is to reconstruct the nullspace policy π from the estimated \mathbf{u}^{ns} . The following describes the two steps in detail.

4.2.1. Step 1—Learning nullspace component. The first step is to extract the nullspace component \mathbf{u}^{ns} from the raw observations \mathbf{x} , \mathbf{u} . As discussed in ref. [16], a requirement on this step is that the data are grouped into multiple subsets such that the constraint matrix \mathbf{A} is consistent within each subset. In the present setting, this separation arises naturally from consideration of the different phases of the gait.

Specifically, since the task constraints might be different across different persons and phases, each phase of each subject is considered an independent subset for the purposes of Step 1. Given input data



as pairs of states and actions $\mathbf{D} = {\mathbf{x}, \mathbf{u}}$, the data-set is divided into *i* subsets such that subset \mathbf{D}_i is the set of observations from subject *i*. Then, each \mathbf{D}_i is further divided into *k* subsets $\mathbf{D}_{i,k} = {\mathbf{x}_{i,k}, \mathbf{u}_{i,k}}$, where $k = {1, 2, 3, ...}$ denotes the phase number.

For the (i, k)th data subset, we seek a model that minimises the inconsistency between the true nullspace component $\mathbf{u}_{i,k}^{ns}$ and the estimated nullspace component $\mathbf{\tilde{u}}_{i,k}^{ns}$.

$$E\left[\mathbf{u}_{i,k}^{ns}, \tilde{\mathbf{u}}_{i,k}^{ns}\right] = \sum_{n=1}^{N_{i,k}} \left\|\mathbf{u}_{i,k,n}^{ns} - \tilde{\mathbf{u}}_{i,k,n}^{ns}\right\|^2.$$
(4)

Since we do not have access to the true nullspace component \mathbf{u}^{ns} , Eq. (4) cannot be directly optimised. Instead, we attempt to eliminate the components of motion that are due to the task constraints, and learn a model that is consistent with the observations. To achieve this, we seek a projection matrix $\mathbf{P}_{i,k} = \frac{\mathbf{\tilde{u}}_{i,k}^{ns} \mathbf{\tilde{u}}_{i,k}^{ns^{\top}}}{\|\mathbf{\tilde{u}}_{i,k}^{ns}\|^2}$, which projects $\mathbf{u}_{i,k}$ onto the learnt nullspace component and satisfies $\mathbf{P}_{i,k} \mathbf{u}_{i,k} \equiv \mathbf{P}_{i,k}(\mathbf{u}_{i,k}^{ts} + \mathbf{u}_{i,k}^{ns}) = \mathbf{\tilde{u}}_{i,k}^{ns}$. The objective function Eq. (4) can be rewritten in terms of this projection as

$$E_1[\tilde{\mathbf{u}}_{i,k}^{ns}] = \sum_{n=1}^{N_{i,k}} \left\| \frac{\tilde{\mathbf{u}}_{i,k,n}^{ns} \left(\tilde{\mathbf{u}}_{i,k,n}^{ns} \right)^{\mathsf{T}}}{\left\| \tilde{\mathbf{u}}_{i,k,n}^{ns} \right\|^2} \mathbf{u}_{i,k,n} - \tilde{\mathbf{u}}_{i,k,n}^{ns} \right\|^2.$$
(5)

Figure 6(a) illustrates an example of this idea. Assuming there are two observations \mathbf{u}_1 , \mathbf{u}_2 from the same \mathbf{x} within a subset. Since the constraint is consistent, \mathbf{u}_1 , \mathbf{u}_2 must have the same \mathbf{u}^{ns} . We seek a \mathbf{u}^{ns} such that when \mathbf{u}_1 and \mathbf{u}_2 are projected onto \mathbf{u}^{ns} , the error is minimised.

In this paper, each $\mathbf{u}_{i,k}^{ns}$ is modelled through iterative optimisation of Eq. (5) using Gaussian radial basis functions. More precisely, $\tilde{\mathbf{u}}_{i,k}^{ns} = \mathbf{W}_{i,k} \boldsymbol{\beta}(\mathbf{x}_{i,k})$, where $\mathbf{W}_{i,k} \in \mathbb{R}^{d \times M}$ is a matrix of weights, and $\boldsymbol{\beta}(\mathbf{x}_{i,k}) = \frac{K(\mathbf{x}_{i,k} - \mathbf{c}_m)}{\sum_{m=1}^{M} K(\mathbf{x} - \mathbf{c}_m)} \in \mathbb{R}^M$ is a vector of basis functions, M is the number of basis functions and \mathbf{c}_m for m = 1, ..., M are the centres. The optimisation is initialised using direct regression to find the initial approximation $\mathbf{W}_{i,k}^0 = \arg \min \sum \|\mathbf{u}_{i,k} - \tilde{\mathbf{u}}_{i,k}^{ns}\|^2$.

4.2.2. Step 2 — Learning nullspace policies. The output of Step 1 is a set of $i \times k$ intermediate models for the nullspace component $\tilde{\mathbf{u}}_{i,k}^{ns} \approx \mathbf{N}_{i,k} \pi$. The goal of the second step is to generalise from these to find an approximate policy $\tilde{\pi}$ that is consistent with all of the estimated nullspace component $\tilde{\mathbf{u}}_{i,k}^{ns}$. Ideally, the approximation should minimise the error between the true policy and the learnt policy:

$$E[\pi, \tilde{\pi}] = \sum_{n=1}^{N} \|\pi_n - \tilde{\pi}_n\|^2.$$
 (6)

Unfortunately, since the true policy π is not observed, Eq. (6) cannot be minimised directly.

Instead, we proceed by noting that, on completion of Step 1 we have the equivalent to a set of $i \times k$ systems that satisfy $\mathbf{A}_{i,k} \tilde{\mathbf{u}}_{i,k}^{ns} = 0$. As a result we can adapt the work in ref. [15] to find a policy that is maximally consistent with the observations.

More precisely, the $i \times k$ intermediate models are combined into a single data-set (\mathbf{x} , $\tilde{\mathbf{u}}^{ns}$), and the approximation is made by minimising the *inconsistency error*

$$E_{2}[\tilde{\boldsymbol{\pi}}] = \sum_{n=1}^{N} \left\| \frac{\tilde{\mathbf{u}}_{n}^{ns} \tilde{\mathbf{u}}_{n}^{ns^{\top}}}{\left\| \tilde{\mathbf{u}}_{n}^{ns} \right\|^{2}} \tilde{\boldsymbol{\pi}}(\mathbf{x}_{n}) - \tilde{\mathbf{u}}_{n}^{ns} \right\|^{2}.$$
(7)

An example is illustrated in Fig. 6(b). Given two (or more) nullspace components \mathbf{u}_1^{ns} and \mathbf{u}_2^{ns} , the inconsistency error favours models for which there is minimal discrepancy between \mathbf{u}_1^{ns} and \mathbf{u}_2^{ns} and the model, projected onto these observations.

In this paper, the nullspace policy $\tilde{\pi}$ is modelled with Gaussian radial basis functions, through minimisation of Eq. (7). The entire process of estimating π for periodic gaits is summarised in Algorithm 1.



Algorithm 1 Nullspace Policy Learning

Input: $D = \{x, u\}$: data-set of states x and action u

Output: $\tilde{\pi}$: learnt null-space policy

- 1: Split **D** into \mathbf{D}_i where \mathbf{D}_i is the input from subject *i*
- 2: Split \mathbf{D}_i into $\mathbf{D}_{i,k}$ where k denotes the phase number
- 3: for all $\mathbf{D}_{i,k}$ do
- 4: Learn $\tilde{\mathbf{u}}_{i,k}^{ns}$ by minimising Eq. (5)
- 5: end for
- 6: Combine { $\mathbf{x}_{i,k}, \tilde{\mathbf{u}}_{i,k}^{ns}$ } into a single data-set { $\mathbf{x}, \tilde{\mathbf{u}}^{ns}$ }
- 7: Learn $\tilde{\pi}$ by minimising Eq. (7)



Fig. 6. (Colour online) A schematic of (a) Step 1 and (b) Step 2 of the proposed learning algorithm.



Fig. 7. (Colour online) Example of different variations in tasks: (a) if only \mathbf{u}_1 is observed, there are more than one way to decompose \mathbf{u}^{ts} and \mathbf{u}^{ns} . (b) By observing \mathbf{u}_1 and \mathbf{u}_2 , \mathbf{u}^{ns} can be determined.

4.3. Feasibility of learning consistent policy

In this section, we discuss the requirements and feasibility of learning a consistent policy across subjects.

4.3.1. Consistent projection matrix in Step 1. The goal of Step 1 is to learn the nullspace components \mathbf{u}^{ns} for estimation of π in Step 2. For this to be effective, \mathbf{u}^{ns} must be consistent within the dataset used for estimation. This implies that π and the projection matrix $(\mathbf{I} - \mathbf{A}^{\dagger} \mathbf{A})$ should both be consistent. Since π is expected to capture the characteristics of walking that are consistent across tasks and embodiments, the main variation in observations is expected to come from variations in the constraints \mathbf{A} .

Variations in A may arise due to variations in embodiments and phases. Consistency of A is ensured, then by breaking the data into independent subsets according to subject and phase.

4.3.2. Variation of tasks in Step 1. The second requirement of Step 1 is that multiple **u** are observed at (or near) the same **x**. This is because if only one state-action pair (**x**, **u**) is observed, there is more than one way to decompose **u** into orthogonal components. $\tilde{\mathbf{u}}^{ns}$ can be anything that satisfies $\tilde{\mathbf{u}}^{ns} \approx \frac{\tilde{\mathbf{u}}^{ns} \tilde{\mathbf{u}}^{ns+1}}{||\tilde{\mathbf{u}}^{ns}||} \mathbf{u}$.

For example, consider the situation in Fig. 7. In Fig. 7(a), since only $\mathbf{u} = [1, 1]^{\top}$ is observed, \mathbf{u}^{ns} can be either $[0, 1]^{\top}$ or $[1, 0]^{\top}$. In Fig. 7(b), if both observations \mathbf{u}_1 and \mathbf{u}_2 are observed, \mathbf{u}^{ns} can be determined.





Table III. Requirements of nullspace policy learning.

Fig. 8. (Colour online) In (a), if only one \mathbf{u}^{ns} is observed, there are more than one solution for π . In (b), observing multiple \mathbf{u}^{ns} determines an unique π . (Figures were modified from ref. [15]).

If the constraint **A** and the policy π are consistent within each subset, the variation of **u** depends on task-space policy **b**. This suggests that, in order to determine a consistent π , data should be collected containing a sufficient variation in tasks, for example by asking subjects to walk at different speeds or with different step-sizes.

4.3.3. Variation of projection in Step 2. The goal of Step 2 is to learn a consistent π given \mathbf{u}^{ns} . A requirement for this is that multiple \mathbf{u}^{ns} are observed at (or near) the same \mathbf{x} .

For example, consider the scenarios depicted in Fig. 8. In Fig. 8(a), if we only have one observation \mathbf{u}^{ns} , there is more than one way to estimate $\boldsymbol{\pi}$. In fact, any vector orthogonal to \mathbf{u}^{ns} can be a solution (e.g., $\boldsymbol{\pi}'$). In contrast, if multiple \mathbf{u}^{ns} are observed (Fig. 8(b)), $\boldsymbol{\pi}$ can be determined.²⁸

In the setting considered in this paper, variations in the projection matrix N arise both from variations in the constraint matrix A across phases, as well as differences in the embodiment of subjects.

The requirements of nullspace policy learning and the ways in which they are fulfilled are summarised in Table III.

4.4. Evaluation

1234

The quality of a policy is evaluated by the following criteria:

4.4.1. Unconstrained policy error. Our primary evaluation criteria is the normalised unconstrained policy error (UPE), which directly compares the true and the learnt policies:

$$UPE = \frac{1}{N\sigma_{\pi}^2} \sum_{n}^{N} \|\boldsymbol{\pi}_n - \tilde{\boldsymbol{\pi}}_n\|^2, \qquad (8)$$

where σ_{π} is the standard deviation (S.D.) of the true policy.

4.4.2. Constrained policy error: In some cases, it may be that the variation in constraints is insufficient to fully uncover the true policy π . However, in such cases, where the constraints exhibit little variation, there may be no need to uncover the hidden components of the fully unconstrained policy (since those components are anyway eliminated by the constraints in normal circumstances). In such



A novel approach for representing and generalising periodic gaits



Fig. 9. Decomposition of simulation data.

circumstances, an alternative quality measure is the normalised constrained policy error (CPE).

$$CPE = \frac{1}{N\sigma_{\pi}^2} \sum_{n}^{N} \left\| \mathbf{u}_n^{ns} - \mathbf{N}_n \, \tilde{\boldsymbol{\pi}}_n \right\|^2.$$
⁽⁹⁾

It measures the difference between the data and the estimated policy, when the latter is projected by the same constraints as in the training data.

4.5. Validation on simulated data

To illustrate the application of the nullspace policy learning approach just described to gait analysis, here we briefly present some numerical results validating its performance. The parameters for these tests (i.e., embodiments, speed, step-size) are chosen so that they are similar to our human data (details will be provided in Section 6.1).

4.5.1. Data. To simulate recordings of walking data from multiple, healthy subjects, data were collected from a set of simulated two-link systems. The subjects were assumed to (i) have different embodiments and (ii) perform different tasks.

Specifically, three two-link systems with different link-lengths were employed, to represent three subjects with different embodiments and each subject (two-link system) was recorded walking at three different speeds ω to represent slow, normal and fast walking. The phase divisions and task constraints were chosen based on the example described in Section 3.2.2. The initial joint-angle \mathbf{x}_0 of each cycle was chosen to match our human data (see Section 6), hence, each gait cycle had a slightly different task **b**. As ground truth nullspace policies, we considered:

- 1. a linear policy: $\pi = \beta^{ns}(\mathbf{x} \mathbf{x}^*)$ where $\mathbf{x}^* = (-90^\circ, -25^\circ)$ was chosen as a 'comfort' position to which the system tends to track.
- 2. a limit-cycle policy: $\dot{r} = r(\rho r^2)$ with radius $\rho = 0.1$, angular velocity $\dot{\theta} = -2$ rad/s, where r and θ were the polar representations of the state (i.e., $\mathbf{x} = (r \cos(\theta), r \sin(\theta))$).
- 3. the same linear policy as (a), but the constraints were imposed in the end-effector space. The constraints were slightly modified to simulate walking on various slopes.

4.5.2. Learning. Figure 9 shows the generative model of the data. The inputs are pairs of \mathbf{x} , \mathbf{u} , and the goal is to recover the policy π .

The nullspace component $\mathbf{u}_{i,k}^{ns}$ was learnt for each phase of each two-link system by minimising Eq. (5), which yields nine different models. Each $\mathbf{u}_{i,k}^{ns}$ was modelled using a set of Gaussian RBFs, where the number of RBFs, N, was obtained through cross-validation for $N \in [10, 50]$. The centres were chosen according to k-means, and the widths σ^2 were the mean distance between centres.

After Step 1, all $(\mathbf{x}_{i,k}, \tilde{\mathbf{u}}_{i,k}^{ns})$ were combined as the input for Step 2. The nullspace policy was modelled separately with another 50 Gaussian RBFs.

In the following, the 10-fold cross-validation results were reported when using 90% of the data-set for training the models and reserving 10% for testing. The performance of our proposed method was evaluated based on the criteria discussed in Section 4.4. We compared our approach with (i) linear regression and (ii) RBF network.

4.5.3. Results. Figure 10 summarises the results of recovering (a) the linear policy, (b) the limit-cycle policy and (c) the linear policy with various slopes.

In all sub-figures, the measurements on the left are the root-mean squared error in joint space (nMSE) and the measurements on the right are the UPE. The error bars are mean \pm S.D. in log scale



Proposed method



• RBF network

Fig. 10. (Colour online) Errors (mean \pm 3 S.D.) of recovering (a) a linear policy, (b) a limit-cycle policy and (c) a linear policy under various environments.



Fig. 11. (Colour online) Motion generated from (a) the true limit-cycle policy, (b) the learnt policy using RBF network and (c) the learnt policy using our proposed method.

over 10 experiments on a hold-out data-set. (Note that S.D. looks larger at the points lower in the plot due to the log scale.)

From this evaluation, we can see that the proposed method is more accurate in predicting the policy both in terms of the nMSE and UPE. Furthermore, since UPE is a direct comparison between the true and the learnt policies, the proposed method is expected to be more accurate under different task constraints and behaviours.

Figure 11 is a visualisation of the true and recovered limit-cycle behaviours over multiple gait cycles. Each gait cycle takes a different speed and step-size. (In the following figures, higher speeds are represented by darker colours and vice versa.) These two rows show the hip angle and the knee angle, respectively. Figure 11(a) shows the motion using the true limit-cycle policy. RBF network (Fig. 11(b)) fails to generate smooth motion, while our method (Fig. 11(c)) reproduces motion that has excellent consistency with the ground truth. Linear regression is worse even than RBF network, so it is omitted from this visualisation.

5. Identifying Pathological Gaits

One of our objectives is to quantify the difference between gaits, and a potential application is gait abnormality detection. The principle is to compare an unknown gait with a reference gait which is expected to be normal and use their difference as the classification criteria.



Linear regression



Fig. 12. Gait abnormality detection by measuring the difference in nullspace policies.

5.1. Quantifying difference between gaits

With our walking-phase model, comparing two gaits is equivalent to comparing two nullspace policies. A general framework is illustrated in Fig. 12; ideally, we want to extract the nullspace policy from the gait of healthy subjects (as the reference policy π^{ref}) and compare to the nullspace policy from a new subject (π^{new}). A difference above a certain threshold would signify a pathological gait.

However, in order to estimate π^{new} and compute this difference, observations from various embodiments are required, which is infeasible for the new subject. Nevertheless, it is still possible to detect abnormalities in gait by comparing the policies *under constraints*. Specifically, instead of measuring the difference between π^{new} and π^{ref} , we can instead evaluate the difference between $\mathbf{N}^{\text{new}} \pi^{\text{new}}$ and \mathbf{N}^{new} are \mathbf{N}^{new} is a projection matrix derived from the constraints of the new subject.

5.1.1. Projection matrix estimation. The projection matrix of the new person \mathbf{N}^{new} is unavailable by assumption, but it is possible to estimate it. The key to our approach comes from the property of a projection matrix. By definition, $\mathbf{u}^{ns} = \mathbf{N}\pi$, where **N** is the projection matrix which projects a vector onto the image space of **N**. Since \mathbf{u}^{ns} is already in the image space of **N**, we must have $\mathbf{N} \mathbf{u}^{ns} = \mathbf{u}^{ns}$.

Based on this insight, an estimate of the projection matrix can be found by searching over the range of possible projections. Recall that $\mathbf{N} = \mathbf{I} - \mathbf{A}^{\dagger} \mathbf{A}$; hence, we need to find an \mathbf{A} that matches the direction of the constraint as closely as possible. In this paper, we have constraints $\mathbf{A} \in \mathbb{R}^{1\times 2}$, so we can model the constraint as a unit vector $\tilde{\mathbf{A}} = [\cos(\theta)\sin(\theta)]$, where $\theta \in [0, \pi]$ covers all possible cases of \mathbf{N} .

To estimate a projection matrix for the *k*th phase of a new person, we seek a θ_k such that the difference between $\mathbf{\tilde{N}}_k^{\text{new}} \mathbf{u}_k^{ns,\text{new}}$ and $\mathbf{u}_k^{ns,\text{new}}$ is minimised. Note that the true nullspace component of the new subject $\mathbf{u}_k^{ns,\text{new}}$ is also unknown, so we use the *estimate of the nullspace component*, instead.

Namely, after learning the nullspace component $\tilde{\mathbf{u}}_{k}^{ns,\text{new}}$ for the *k*th phase, we obtained N_{k} pairs of $(\mathbf{x}_{k}, \tilde{\mathbf{u}}_{k}^{ns,\text{new}})$. For each possible value of $\theta_{k} \in [0, \pi]$, we can calculate the difference between $\tilde{\mathbf{N}}_{k}^{\text{new}} \tilde{\mathbf{u}}_{k}^{ns,\text{new}}$ and $\tilde{\mathbf{u}}_{k}^{ns,\text{new}}$:

$$E_{\theta}\left[\tilde{\mathbf{N}}_{k}^{\text{new}}\right] = \sum_{n}^{N_{k}} \left\|\tilde{\mathbf{N}}_{k}^{\text{new}}\tilde{\mathbf{u}}_{k,n}^{ns,\text{new}} - \tilde{\mathbf{u}}_{k,n}^{ns,\text{new}}\right\|^{2}.$$
(10)

The optimal $\theta_k \in [0, \pi]$ minimises Eq. (10).

5.1.2. Approximate policy difference. Given the estimated reference policy π^{ref} from the healthy subjects and measurements of the behaviour of the new person, we can now quantify the difference of the new person's policy through a measure that we call the *approximate constrained policy difference* (APD).

Specifically, given pairs of state/action observations of the new person \mathbf{x}^{new} , \mathbf{u}^{new} , we first divide their data into *K* walking phases and learn a model of the nullspace components $\tilde{\mathbf{u}}_{k}^{ns,\text{new}}$ for each phase. Then, we take $\tilde{\mathbf{u}}_{k}^{ns,\text{new}}$ and estimate the projection matrix $\tilde{\mathbf{N}}_{k}^{\text{new}}$ using Eq. (10). The APD is computed as

$$APD = \frac{1}{N\sigma_{\pi^{\text{ref}}}^2} \sum_{k=1}^K \sum_{n=1}^{N_k} \left\| \tilde{\mathbf{u}}_{k,n}^{ns,\text{new}} - \tilde{\mathbf{N}}_k^{\text{new}} \,\boldsymbol{\pi}_{k,n}^{\text{ref}} \right\|^2, \tag{11}$$



System	Nullspace policy	Task constraints	
S1	Linear	Same as reference	
S2	Linear	Different from reference	
S 3	Sinusoid	Same as reference	
S4	Sinusoid	Different from reference	

Table IV. New systems for testing. The reference is the policy learnt from the linear data-set in Section 4.5.

where $\sigma_{\pi^{ref}}^2$ is the variance of the reference policy. The APD measures the difference between the reference policy and the new person in the constrained space, normalised by the variance of the reference policy. Algorithm 2 summarises the process.

Algorithm 2 Approximate difference between two gaits
Input: $\mathbf{D} = {\mathbf{x}^{new}, \mathbf{u}^{new}}$: data-set of a new person
π^{ref} : reference policy
Output: APD: approximated policy difference
1: Split D into \mathbf{D}_k where k denotes the phase number
2: for all \mathbf{D}_k do
3: Learn $\tilde{\mathbf{u}}_k^{ns,\text{new}}$ by minimising Eq. (5)
4: Learn $\tilde{\mathbf{N}}_k^{\text{new}}$ by minimising Eq. (10)
5: end for
6: Approximate the difference using Eq. (11)

One way to interpret the quantity in Eq. (11) is, this is the difference between a person and the reference in an abstract policy space, where the distance resulting from various walking behaviours (speeds, step-sizes) is eliminated. Therefore, this measurement can be thought of as the quantification of how much we should correct a gait without interfering with his/her speed or step-length. Additionally, the APD can also be used for monitoring a program of rehabilitation by looking at the evolution of the APD with training.

5.2. Experiments

1238

In this section, we demonstrate this idea using artificial data. Imagine a scenario in which there are four previously unseen persons, two of whom have healthy gait and the other two have pathological gait. The goal is to assess the use of the APD in quantifying the difference between the behaviour of each person and healthy walking, as captured by the learnt reference policy.

5.2.1. Linear policy as reference policy. In our first evaluation, we used the learnt linear policy from Section 4.5 as the reference policy (i.e., π^{ref} in Fig. 12), and compared this against data from four additional two-link systems representing four previously unseen subjects (we will refer to them as S1, S2, S3, S4).

Among these four subjects, we used S1 and S2 to represent the normal gait, i.e., followed the same linear policy as the reference. We used S3 and S4 to represent the abnormal gait; namely, they followed a different (sinusoidal) policy $\pi^{\sin}(\mathbf{x}) = -\sin(\mathbf{x}^* - \mathbf{x})$, where $\mathbf{x}^* = (-90^\circ, -25^\circ)$ was the nullspace target.

The new subjects took the same phase divisions and task-space policies as described in Section 3.3. To explore the effect of different task constraints, S1 and S3 had the same task constraints with one of subjects in the reference data, while S2 and S4 had a totally different one. Note that the nullspace policy is independent of the task constraints, and therefore, the results are expected to be consistent even when the task constraints change. The setup of these four systems are listed in Table IV.

Algorithm 2 was applied on each new system separately. More specifically, we learnt the $\mathbf{u}^{ns,new}$ and \mathbf{N}^{new} for each phase of each person. The differences between each system and the reference were calculated using APD (Eq. (11)). We also tested linear regression and RBF network for comparison.



System	Nullspace policy	Task constraints
S 1	Limit-cycle	Same as reference
S2	Limit-cycle	Different from reference
S 3	Sinusoid	Same as reference
S4	Sinusoid	Different from reference

Table V. New systems for testing. The reference is the policy learnt from the limit-cycle data-set in Section 4.5.



Fig. 13. (Colour online) Average difference between the testing systems and the reference, where the reference is (a) a linear policy and (b) a limit-cycle policy. The error-bars are mean \pm S.D. in log-scale over 10 experiments. The results were grouped into normal (N) and abnormal (A), where N is the average of S1 and S2, and A is the average of S3 and S4.

We took the learnt model in Section 4.5, and measured the difference in joint-space between the reference policy and each new system.

Figure 13 summarises the results of comparing the reference policy to S1–S4 using three different methods. The *y*-axis shows the average differences in joint-space for the regression methods and the average APD for the proposed method. (Note that the proposed method attempts to eliminate the difference resulting from various walking behaviours; therefore, our method does not directly compare joint-angles, which is influenced by various behaviours.) The error bars are the mean \pm S.D. in log scale over 10 experiments.

In Fig. 13(a), the horizontal line is the true differences between the reference (linear) and S3–S4 (sinusoidal), which is the true difference between $N\pi^{sin}$ and $N\pi^{ref}$. Note that S1 and S2 adapt the same linear policy as the reference, and the true differences between S1, S2 and the reference are 0.

In Fig. 13(a), the yellow and green colours denote the results of standard methods. We can see that linear regression fails to differentiate linear and sinusoidal policies. Although RBF network predicts relatively higher difference for abnormal (A), the predicted difference for normal (N) is also unreasonably high. The red colour denotes the results of our approach. Our method yields lower error for normal (N), and the result confirms with the fact that S1 and S2 use the same linear policy. The error for abnormal (A) is relatively higher, which is also expected since S3 and S4 adopt the sinusoidal policy.

5.2.2. *Limit-cycle policy as reference policy*. After validating with a linear policy, we also tested our method on a non-linear policy. We took the limit-cycle policy learnt in Section 4.5 as the reference policy. S1 and S2 were generated using the same limit-cycle policy as the reference (see Table V).

In Fig. 13(b), the horizontal line shows the true difference between the reference policy (limit-cycle) and S3–S4 (sinusoidal), which is the true distance between N π^{sin} and N π^{ref} .

Similar to the last experiment, the standard methods (yellow and green) fail to make reasonable predictions. Our method (red), again, produces results relatively similar to the ground truth. Our



Subject	Upper leg (cm)	Lower leg (cm)
S 1	37.7	38.9
S2	41.6	41.1
S 3	42.6	40.5
S4	44.2	43.3
S5	45.1	44.2
S6	42.5	41.5
S 7	45.8	44.8

Table VI. Leg lengths of the subjects.



(b)

Fig. 14. (Colour online) (a) Xsens motion capture system. (b) 3.5-kg weight was attached to the subject's leg to create 'abnormal' gaits.

experiment demonstrates good assessments on quantifying the difference between policies, even if the true policy, the constraints and the tasks are unknown.

6. Experiment on Human Data

After validating our approach on simulated data, in this section, we explored the utility of our approach on a more realistic setting. We collected human kinematic data using motion-capture system and tested how well we can employ the learnt policy for gait abnormality detection.

6.1. Data collection

The kinematics data were collected using Xsens MVN BIOMECH system.²⁹ The sensor units were attached to the subjects according to Xsens configuration (Fig. 14(a)). The update frequency is set to 120 FPS.

6.1.1. Variations. The data were collected from seven males, aged between 20–29 years (referred as S1–S7). These seven subjects have different body types and they were chosen to ensure our data contain some variations in embodiments. The leg-lengths of subjects are summarised in Table VI.

Data were recorded for five walking speeds: 93, 106, 119, 129, 140 steps per minute, which were taken from the speed range reported in ref. [30]. The walking speeds were controlled through the use of a metronome. The subjects were asked to walk such that heel strike coincided with the tick of the metronome. For each speed, 10 walking trials were collected.

6.1.2. Pathological gait. To create 'pathological' gait, 3.5-kg bags of sand were strapped to the subjects' left leg (see Fig. 14(b)). We used the same setup (speed, number of trails, etc.) to collect kinematic data of abnormal gait from each subject.

6.1.3. Pre-processing. The captured motion was processed using MVN Studio 3.0 (a graphical interface provided by Xsens). The data were exported in MVNX format, which is an XML format with full kinematics of each segment, including position, velocity, acceleration, orientation, angular





Fig. 15. (Colour online) The normal gait (red) and the abnormal gait (black) from one subject.

velocity and angular acceleration. The desired parameters were calculated and preprocessed using MATLAB.

We used heel-strike of the left leg to demark the beginning of a walking cycle and extracted as many cycles as possible from all walking trails we collected. We obtained roughly 200 gait cycles from each subject, and 100 gait cycles were selected for analysis. In this experiment, we tested our approach with three walking phases according to the descriptions in Section 3.3.

Figure 15 shows the hip and knee angles from one of the subjects. One trajectory from each walking speed was selected. Note that normal gaits (red) and abnormal gaits (black) look very similar from direct observation.

6.2. Setup

6.2.1. Learning reference policy. We selected five subjects, S1, S2, S3, S4, S5 for learning the reference policy (Algorithm 2). For each walking phase of each subject, we learnt a model for nullspace component $\tilde{\mathbf{u}}_{i,k}^{ns,\text{ref}}$, which yielded 15 models. Each $\tilde{\mathbf{u}}_{i,k}^{ns,\text{ref}}$ was consisted of *M* Gaussian RBFs where *M* varied from 16 to 100. The nullspace policy π^{ref} was also modelled as parametric model with Gaussian RBFs. We used this learnt policy as the reference policy for gait abnormality detection.

6.2.2. Identify pathological gaits. Five subjects (S1–S5) were used to collect five normal and five pathological gaits (using the leg loading). To investigate how well the learnt policy can generalise across subjects, we also performed the same experiment on the subjects whose data had not been used for training the reference policy (S6–S7)—we collected normal and pathological gaits for each of these.

6.2.3. Baseline. For comparison, we also trained models using (i) linear regression and (ii) RBF network on raw observations (\mathbf{x}, \mathbf{u}) from the normal gaits, and tested if we can see a difference between normal and abnormal gaits in joint space.

6.3. Results

Figure 16 shows the average results over all subjects. The yellow and green colours denote the results using standard methods, and the red colour denotes the results using our method. The error bars are the mean \pm S.D. in log scale over 10 experiments.

Figure 16(a) shows the average results over S1–S5. We can see that the standard methods cannot differentiate normal and abnormal gait. Our approach achieved relatively lower difference when comparing with normal gaits and higher difference when comparing with abnormal gaits. Even if we have no access to the true policy, constraints, nor tasks, our reference policy is more effective in differentiating normal from abnormal.

Figure 16(b) shows the average results over S6 and S7. In this case, RBF network predicts that the abnormal gait is more similar to the reference gait. This outcome reflects the problem of using average template, where the reference gait fails to adapt to the new subjects. Our proposed method, on the other hand, can still show some difference between normal and abnormal, even if S6 and S7 are different from the subjects used to train the reference policy.



• Linear regression • RBF

• RBF network • Proposed method



Fig. 16. (Colour online) Average difference between the testing subjects and the reference policy, where the testing subjects are (a) S1–S5 and (b) S6–S7. The error-bars are mean \pm S.D. in log scale over ten experiments and over subjects. The results are grouped into normal (N) and abnormal (A).

RBF network
 Proposed method



Fig. 17. (Colour online) Average difference between the new subjects (S6 and S7) and the reference gait (learnt from S1–S5). The testing data were divided by three different speeds: (a) slow, (b) average and (c) fast.

6.3.1. Apply the learnt policies on various behaviours. We also evaluated how well the learnt policies can generalise across various walking behaviours. In this experiment, we tested slow, average and fast walks separately, and see how the results might be affected by various speeds.

Similar to the last experiment, we used the normal and abnormal gaits from S6 and S7 to represent two normal and two abnormal persons, and compare them with the reference policy learnt from S1, S2, S3, S4 and S5. Figure 17 shows the results of predicted difference between the new subjects and the reference policy, where the data were divided by three different speeds: (a) slow, (b) average and (c) fast.

From Fig. 17, we can see that, by using RBF network (green), the predicted difference between the normal gait and the reference increases as the walking speed increases. This is equivalent to considering those faster walks are deviations from the normal gait.

In contrast, our proposed method (red) yields consistent results regardless of walking speeds, and this outcome confirms the fact that our method attempts to eliminate the difference coming from various walking speeds. Implication in real world application is that our quantification method deals with different walking behaviours consistently, and the patients can choose to walk faster or slower.



7. Conclusion

We explore the problem of representing, generalising and comparing gaits. We consider that locomotion can be described as a combination of characteristics of the gait and variations from embodiments and behaviours. We assume that there exist some consistent characteristics across different embodiments and walking phases, and we aim to generalise them.

Our approach is built upon the work in the domain of nullspace policy recovery and gait analysis. We formulate our problem into a walking-phase model, and we adapt the nullspace policy learning method to generalise a policy that can capture the consistent characteristics of walking gait. Our experiment has shown that our method is effective in reconstructing the policy, even if the true policy, the constraints and the variations in behaviours are unknown. After recovering the policy, we can utilise this policy for gait abnormality detection.

For simplicity, we focus on the kinematics of movement; however, with some modification the model may also be extended to the case of redundancy in dynamics. Future work will focus on exploring different representations and higher degrees of freedom on human data.

One of the current issues in robot-assisted rehabilitation is that the existing systems normally restrict the patients to walk with some pre-defined reference trajectories such that the patients are required to walk at a certain speed, slope or step-size. Our approach is different from the previous work in that we aim to extract the important aspect of normal gaits and separate those from variations. In the future, we would like to use the reconstructed policy to produce the reference trajectory for gait rehabilitation. With our method, the system will be able to correct the patient only if there is a mismatch in the characteristics of walking without interfering his/her personal preferences such as speeds or step-lengths.

References

- F. Multon, L. France, M.-P. Cani-Gascuel and G. Debunne, "Computer animation of human walking: A survey," J. Vis. Comput. Animat. 10(1), 39–54 (1999).
- R. Boulic, N. Thalmann and D. Thalmann, "A global human walking model with real-time kinematic personification," Vis. Comput. 6, 344–358 (1990).
- 3. D. Pongas, M. Mistry and S. Schaal, "A Robust Quadruped Walking Gait for Traversing Rough Terrain," *Proceedings of the International Conference on Robotics and Automation*, Rome, Italy (Apr. 10–14, 2007) pp. 1474–1479.
- 4. Y. Fukuoka, H. Kimura and A. H. Cohen, "Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts," *Int. J. Robot. Res.* **22**(3–4), 187–202 (2003).
- A. Duschau-Wicke, J. von Zitzewitz, A. Caprez, L. Lunenburger and R. Riener, "Path control: A method for patient-cooperative robot-aided gait rehabilitation sky," *IEEE Trans. Neural Syst. Rehabil. Eng.* 18(1), 38–48 (2010).
- P. Artemiadis and H. Krebs, "On the Potential Field-Based Control of the Mit-Skywalker," *Proceedings of the International Conference on Robotics and Automation*, Shanghai, China (May 9–13, 2011) pp. 1427–1432.
- 7. J. Veneman, R. Kruidhof, E. Hekman, R. Ekkelenkamp, E. van Asseldonk and H. van der Kooij, "Design and evaluation of the lopes exoskeleton robot for interactive gait rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.* **15**, 379–386 (2007).
- 8. J. Hidler, D. Nichols, M. Pelliccio and K. Brady, "Advances in the understanding and treatment of stroke impairment using robotic devices," *Top. Stroke Rehabil.* **12**, 22–35 (2005).
- 9. S. Jezernik, G. Colombo and M. Morari, "Automatic gait-pattern adaptation algorithms for rehabilitation with a 4-DOF robotic orthosis," *IEEE Trans. Robot. Autom.* **20**(3), 574–582 (2004).
- 10. J. F. Veneman, R. Ekkelenkamp, R. Kruidhof, F. C. Van Der Helm and H. Van Der Kooij, "A series elasticand bowden-cable-based actuation system for use as torque actuator in exoskeleton-type robots," *Int. J. Robot. Res.* **25**, 261–281 (2006).
- R. Riener, L. Lunenburger, S. Jezernik, M. Anderschitz, G. Colombo and V. Dietz, "Patient-cooperative strategies for robot-aided treadmill training: first experimental results," *IEEE Trans. Neural Syst. Rehabil. Eng.* 13(3), 380–394 (2005).
- 12. V. Stokes, C. Andersson and H. Forssberg, "Rotational and translational movement features of the pelvis and thorax during adult human locomotion," *J. Biomech.* **22**(1), 43–50 (1989).
- 13. Y. P. Ivanenko, R. E. Poppele and F. Lacquaniti, "Five basic muscle activation patterns account for muscle activity during human locomotion," *J. Physiol.* **556**, 267–282 (2004).
- 14. I. T. Jolliffe, Principal Components Analysis (Springer-Verlag, 1986).
- 15. M. Howard, S. Klanke, M. Gienger, C. Goerick and S. Vijayakumar, "A novel method for learning policies from variable constraint data," *Auton. Robots* **27**, 105–121 (2009).
- 16. C. Towell, M. Howard and S. Vijayakumar, "Learning Nullspace Policies," *Proceedings of the International Conference on Intelligent Robots and Systems*, Taipei, Taiwan (Oct. 18–22, 2010) pp. 241–248.



- 17. S. Lee and Y. Sankai, "Power Assist Control for Walking Aid with HAL-3 based on Emg and Impedance Adjustment Around Knee Joint," *Proceedings of the International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland (Sep. 30–Oct. 5, 2002) pp. 1499–1504.
- 18. M. W. Whittle, Gait Analysis: An Introduction, 3rd ed. (Butterworth-Heinemann, Oxford, UK, 2003).
- Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi and H. Hirukawa, "Biped Walking Pattern Generation by Using Preview Control of Zero-Moment Point," *Proceedings of the International Conference on Robotics and Automation*, Taipei, Taiwan (Sep. 14–19, 2003) pp. 1620–1626.
- Y. Ogura, H. Aikawa, K. Shimomura, H. Kondo, A. Morishima, H. Lim and A. Takanishi, "Development of a New Humanoid Robot Wabian-2," *Proceedings of the International Conference on Robotics and Automation*, Orlando, USA (May 15–19, 2006) pp. 76–81.
- 21. Y. Guan, E. S. Neo, K. Yokoi and K. Tanie, "Stepping over obstacles with humanoid robots," *IEEE Trans. Robot.* **22**, 958–973 (2006).
- A. J. Ijspeert, J. Nakanishi and S. Schaal, "Learning Attractor Landscapes for Learning Motor Primitives," In: Advances in Neural Information Processing Systems (S. Becker, S. Thrun and K. Obermayer, eds.) (MIT Press, Cambridge, 2003) pp. 1547–1554.
- J. Morimoto and C. Atkeson, "Nonparametric representation of an approximated poincaré map for learning biped locomotion," *Auton. Robots* 27, 131–144 (2009).
- A. Ijspeert and A. Crespi, "Online Trajectory Generation in an Amphibious Snake Robot Using a Lamprey-Like Central Pattern Generator Model," *Proceedings of the International Conference on Robotics and Automation*, Roma, Italy (Apr. 10–14, 2007) pp. 262–268.
- Y. Nakamura, T. Mori, M. aki Sato and S. Ishii, "Reinforcement learning for a biped robot based on a cpg-actor-critic method," *Neural Netw.* 20(6), 723–735 (2007).
- 26. D. A. Winter, "Kinematic and kinetic patterns in human gait: Variability and compensating effects," *Hum. Mov. Sci.* **3**(1–2), 51–76 (1984).
- D. Broomhead and D. Lowe, "Radial basis functions, multi-variable functional interpolation and adaptive networks," *Complex Syst.* 2, 321–355 (1988).
- M. Howard and S. Vijayakumar, "Reconstructing Null-Space Policies Subject to Dynamic Task Constraints in Redundant Manipulators," *Proceedings of the Workshop on Robotics and Mathematics*, Coimbra, Portugal (Sep. 17–19, 2007) pp. 109–115.
- D. Roetenberg, H. Luinge and P. Slycke, "Xsens mvn: Full 6-DOF Human Motion Tracking Using Miniature Inertial Sensors," Xsens Technologies BV, Enschede, The Netherlands, 2009.
- T. Öberg, A. Karsznia and K. Öberg, "Basic gait parameters: Reference data for normal subjects, 10–79 years of age," J. Rehabil. Res. Dev. 30, 210–210 (1993).

