Online ISSN 1848-3380 , Print ISSN 0005-1144
ATKAFF 52(3), 223–232(2011)

Sven Wehner, Maren Bennewitz

Humanoid Gait Optimization Based on Human Data

UDK 004.896.02
IFAC 2.6; 5.9.3

Original scientific paper

Achieving a stable, human-like gait for humanoid robots is a challenging task. While a variety of techniques exist to generate stable walking patterns, only little attention has been paid to the resemblance to the human gait. Popular gaits, for example, apply the strategy to bend the knees and to swing the torso in the lateral direction in order to ensure stability by shifting the center of mass. As a result, the walking patterns do not look very human-like. However, human resemblance is an important aspect whenever robots are designed to coexist and interact with humans. In this article, we present techniques to optimize a given, stable gait of a humanoid robot with respect to human resemblance. To acquire human data, we use a full-body motion capture system. We propose four different optimization algorithms that work at joint angle basis and use the joint angle difference as measure of similarity. The experiments carried out with a *HOAP-2* robot in simulation demonstrate that all techniques generate a gait that is significantly more human-like compared to the robot's initial gate. As the results show, the optimization methods based on hill climbing and policy gradient estimation yield the best performance.

Key words: Humanoid robots, Gait optimization, Human-like walking

Optimizacija držanja čovjekolikih robota utemeljena na podacima iz čovjekovog hoda. Postizanje stabilnog, čovjekolikog držanja čovjekolikih robota vrlo je zahtjevan zadatak. Iako postoji mnoštvo tehnika koje se koriste za postizanje stabilnih uzoraka hodanja, malo se pažnje pridaje sličnosti s ljudskim držanjem. Primjerice, neke tehnike koriste strategiju savijanja koljena i njihanja torza u lateralnom smjeru kako bi se osigurala stabilnost kroz promjenu položaja centra mase. Kao rezultat toga, uzorci hodanja nisu slični čovjekovom hodu. Ipak, sličnost čovjeku važan je aspekt u slučajevima kada su roboti izvedeni za suživot i interakciju s ljudima. U ovom su članku predstavljene tehnike za optimizaciju danih stabilnih držanja čovjekolikog robota s ciljem sličnosti čovjeku. Za prikupljanje podataka o čovjeku koristi se sustav za snimanje cjelokupnog gibanja tijela. Predložena su četiri različita algoritma optimizacije koji koriste zajedničku derivaciju kuta kao mjeru sličnosti. Eksperimenti provedeni koristeći HOAP-2 robot pokazuju da svi postupci generiraju držanje koje je značajno više nalik čovjekovom držanju u odnosu na početno držanje robota. Kao što rezultati pokazuju, metode optimizacije koje se zasnivaju na penjanju uzbrdo i estimaciji gradijenta smjera daju najbolje rezultate.

Ključne riječi: čovjekoliki roboti, optimizacija držanja, čovjekoliko hodanje

1 INTRODUCTION

Recently, humanoid robots have been enjoying great popularity and are now used as a research tool in many groups worldwide. These types of robots possess human-like actuators and sensors that allow them to act in environments designed for humans. One of the motivations behind the research area of humanoid robotics is to develop robots that are able to coexist with humans and interact with them in a natural way.

Compared to wheeled robots, one challenge when working with humanoid robots is to design stable walking gaits for their biped locomotion. In this context, a common technique is to apply heuristics and to manually configure walking patterns for humanoids and carefully choose their parameters. For example, there exist techniques based on central pattern generators (CPGs) to generate joint trajectories using nonlinear oscillators [1, 2]. In these approaches, it is a challenging problem to find appropriate parameters to achieve a stable gait. More computational demanding methods use the concept of the zero moment point (ZMP) [3] and rely on joint angle trajectories, which are computed considering dynamic motion of the robot. Here, an accurate model of the robot and its dynamics is needed to ensure reliable execution.

Several approaches have been presented that aim at optimizing properties such as speed [4–6] or torso stability [7, 8] of a humanoid's walk. The resulting, optimized walking patterns often look rather unnatural. However, the resemblance to the human ideal should also be taken into account when generating walking patterns for humanoid

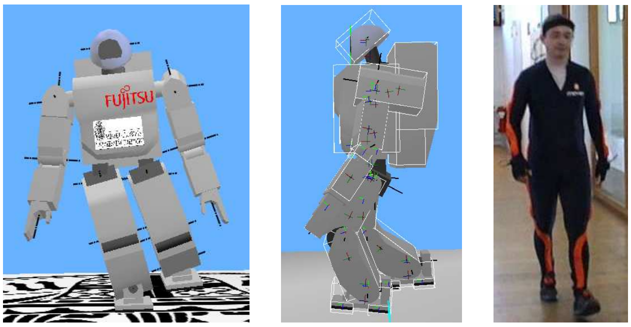

Fig. 1. The left and middle image show snapshots of the robot's initial gait. As can be seen, the robot's torso moves extremely in the lateral direction and the knees are unnaturally bent. In the right image, a human wearing the suit of the full-body motion capture system MVN is walking during the acquisition of the human data.

robots. Since these robots are designed to coexist and interact with humans, it is important that their motions look human-like to guarantee a broad acceptance.

The first two images of Fig. 1 show snapshots of a typical humanoid's gait. As can be seen, the robot shifts its center of mass during walking to ensure stability. First, the torso swings extremely in the lateral direction and, second, the knees are unnaturally bent. These two concepts are often used in humanoid locomotion. However, the resulting gait does not look very natural and human-like.

In this article, we consider the problem of achieving a stable, human-like gait with a humanoid robot. We treat this as an optimization problem and develop four algorithms that work on joint angle basis. Our optimization starts from an initial, stable gait of the robot obtained via a CPG. We consider the walk of a human recorded by a motion capture system (see right image of Fig. 1) as the ideal gait the optimization is aiming at. We define the similarity between the gaits in terms of joint angle difference between the human's and the robot's joint angle trajectories. The advantage of working solely on the basis of joint angles is that we do not need to incorporate expert knowledge into the learning process, e.g., in form of a parameterized gait [7] or in form of a segmentation into different walking phases [9].

We extensively evaluated and empirically compared the different optimization techniques in experiments carried out in the *Webots* simulator [10] with a *HOAP-2* [11] robot. The results show that the optimization methods based on hill climbing and on policy gradient estimation perform best. However, all techniques are able to improve the initial gait so that it is significantly more similar to the human gait.

The article is organized as follow. After discussing related work in the next section, we present the humanoid robot used for the experiments in Sec. 3 and describe the collection and preprocessing of human data in Sec. 4. In Sec. 5, we introduce the algorithms we developed to optimize the robot's walk. Finally, we discuss the experimental results in Sec. 6.

2 RELATED WORK

In the last few years, several techniques to optimize the behavior or the motions of humanoid robots have been presented.

The vast majority of these approaches optimize the humanoid's walking speed for a given, parameterized gait. For example, Faber and Behnke [4] applied an optimization based on policy gradient reinforcement learning and particle swarm optimization to increase the forward speed. The authors used eight parameters of the gait and developed two feedback mechanisms that were included into the optimization process. Niehaus *et al.* [5] also applied particle swarm optimization to speed up the walking capabilities of a humanoid. They considered 14 parameters and performed an optimization of the parameters for different walking directions to allow for omnidirectional walking. Hemker *et al.* [6] applied sequential surrogate optimization to searching for optimal values of the five chosen gait parameters of forward walking. Geng [12] *et al.* proposed a policy gradient reinforcement learning approach to optimize the parameters of a neuronal sensor-driven controller for a planar biped. Furthermore, several researches applied machine learning techniques to optimize the gait of quadruped robots (e.g. [13, 14]). Note that stability is not such a serious problem with quadruped robots, which makes optimization easier. In contrast to all techniques presented above, our goal is not to optimize the speed of a humanoid robot, which often leads to rather unnatural looking walking behaviors. Instead, we aim at generating a gait that looks more human-like than the initial walking behavior of the robot.

Chalodhorn *et al.* [8] used an imitation-based approach to teach a humanoid robot stable walking. The authors recorded human data with an optical motion capture system. They proposed to search for appropriate actions leading to a stable robot gait in a dimensionality-reduced space of the joint angles. The authors enforced stability by using gyroscope signals to favor upright torso positions. From the presented images, it seems that the human demonstrator moved in a rather unnatural way in order to facilitate the learning of the robot. Huang *et al.* [7] analyzed the characteristics of the human gait in terms of change of a number of manually defined walking parameters when changing the step length and the walking cycle. Afterwards, they considered the learned characteristics and determined parameters for a walking pattern, which resulted in a high upper-body stability of a humanoid robot. The authors also used human motion data captured by an optical system. Serhan *et al.* [9] proposed to extract "critical angles" from human locomotion that influence the speed and step length. They used a segmentation of the walking cycle in eight phases and defined maximum angles for the different phases. The authors presented experiments in which a simulated biped robot with a 4 DOF trunk achieved a dynamic walk. To generate a human-like walking behavior that can be adapted according to observations, e.g., barriers or stairs, Denk and Schmidt [15] proposed to concatenate previously learned walking primitives.

Regarding other tasks than walking, several researchers have concentrated on generating motions for humanoids that are optimal with respect to specific criteria. For example, Harada *et al.* [16] optimized motion primitives for a humanoid robot in terms of joint torque, acceleration, and angular momentum. The authors identified relevant variables for the individual tasks (e.g., a reaching motion, walking, or climbing). Faber *et al.* [17] proposed a strategy for body and gaze control of a humanoid during human-robot interaction. They considered different factors, i.e., tracking error, discomfort, and effort to control the pitch and yaw joints to generate human-like turning motions of the robot. Bobrow *et al.* [18] optimized the motion of robots performing different tasks with respect to minimum control effort. Svinin *et al.* [19] considered the problem of generating human-like reaching movements with a robotic arm. The authors use an objective function that is based on the minimization of hand jerk.

In the remainder of this article, we present our approach to optimizing the walking motion of humanoids with respect to human resemblance. Our optimization works on joint angle basis and does not need to incorporate expert knowledge, e.g., for defining appropriate parameters.

3 THE HUMANOID ROBOT

For our experiments, we use a simulated *HOAP-2* from *Fujitsu* [11]. The first two images of Fig. 1 show the simulated robot while walking. The robot has a weight of 7 kg and is 50 cm tall. The total number of its degrees of freedom (DOFs) is 25, but we only consider 21 as relevant for walking, i.e., we excluded head and hand joints. Fig. 2 gives a detailed overview over the robot's degrees of freedom.

We use the model of the *HOAP-2* for the *Webots* simulator. The walking motion, which we take as a basis for our experiments, was generated using a central pattern generator. In this gait, the robot shows an extreme lateral swing and walks with strongly bent knees (see Fig. 1).

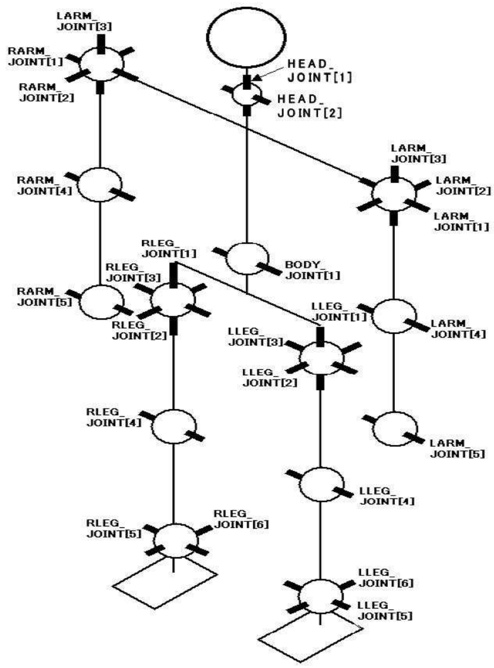

Fig. 2. Overview over the robot's joints. There are six degrees of freedom in each leg, five in each arm, two in the head and one in the hip. The bold lines represent the rotation axes of the joints. Illustration taken from [11].

There are obviously some differences between the human's and the robot's anatomy which can lead to different behavior. One of the major issues is the fact that the robot wears a "backpack" in which the processing unit is located. Accordingly, the robot's center of mass is shifted to the back. Furthermore, certain degrees of freedom of the human body are not represented by a joint in the robot's model, as detailed in the next section.

The given simulation model provides only a single pressure sensor in each foot. We extended the simulated model so that it has a sensor in the front and another one in the back of each foot. We need this extension for our approach of stability estimation.

4 DATA ACQUISITION AND PREPROCESSING

In this section, we describe how to acquire, process, and transform human data into the joint angle trajectories, which are finally used as input to our optimization framework.

4.1 Recording Human Gait Data

For recording full-body motions of a human, we use the *Xsens MVN* [20] system. This motion capture device consists of sixteen sensors, which include accelerometers, gyroscopes, and magnetometers. The right picture

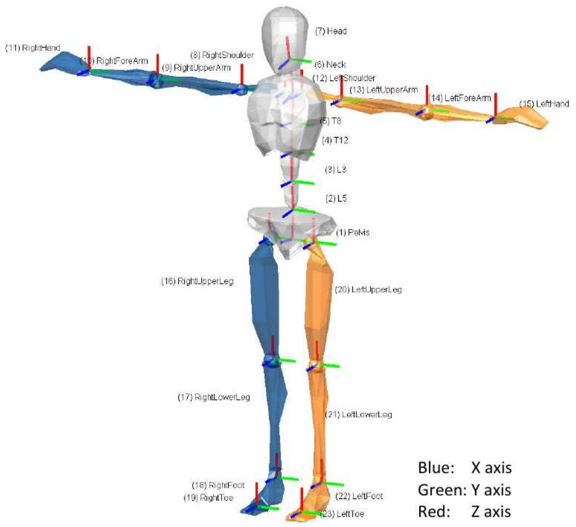

Fig. 3. Position of the various segments which are covered by motion capture system MVN. Illustration taken from [20].

in Fig. 1 shows the human demonstrator wearing the *MVN* body suit. The system records the data with a frequency of 100 Hz. The output includes information about the position and rotation of the 23 body segments. Fig. 3 illustrates the positions of the segments used by the system.

Since we use a joint angle based representation, we transform the rotation information into joint angles. Most mappings of human joint angles to the robot's ones are straightforward. We simply ignore joint angles not existing in the robot, e.g., the rotation around the torso's yaw and roll axes. Furthermore, we need to approximate a small number of robot joints using the human data. For example, we compute the shoulder joint angles using the upper arm and neck orientation.

4.2 Data Preprocessing

After recording the data, we perform some data preprocessing steps in order to get the input data for our optimization framework. This preprocessing is necessary to deal with noise in the captured motion data and with irregularities in the human's gait and to temporally align the human's and the robot's trajectories.

4.2.1 Segmentation of Sequences

We first extract a sequence of five gait cycles (five double-steps) of the human's and of the robot's gait which we use for analysis and on which we perform the optimization. Our experiments indicated that a number of five double-steps was sufficient. The extrema in the robot's and the human's knee trajectories are used to determine the start and end of a sequence.

4.2.2 Noisy Data

Obviously, the captured human data is noisy to a certain extend and the human's motions are not always exactly the same in each gait cycle. To reduce the influence of these effects, we compute the average over different sequences containing five human gait cycles. In particular, we use *Dynamic Time Warping* (*DTW*) [21]. The main idea of DTW is to match two sequences by warping the temporal position of data points in a way that the lowest overall distance between both is found. The distance of two joint angle trajectories $\boldsymbol{M}$ and $\boldsymbol{N}$ is defined as

$$\mathrm{dist}\,(\boldsymbol{M},\boldsymbol{N}) = \sum_{\text{joint } j}\sum_{\text{time } t}\left(M_{j,t}-N_{j,t}\right)^2 . \qquad (1)$$

The DTW algorithm finds the optimal matching between data points and minimizes the overall distance of the sequences. Afterwards, the average over the two sequences is computed by using data points corresponding to the same time steps. We perform DTW in a tree-like manner and subsequently match a sequence with the result of the previous DTW process.

4.2.3 Symmetry

In the recorded data, the trajectories of corresponding lateral joints (e.g., the left and the right knee) of the human appeared to slightly differ. Such small asymmetries can result in instable behavior of the robot. We therefore compute the joint angle trajectories as the average over the data points of corresponding left and right joints.

4.2.4 Uniform Gait Cycles

Afterwards, we eliminate irregularities within the different gait cycles of the captured human trajectories. This includes different execution times of steps or also slightly varying trajectories. We unify the trajectories corresponding to different gait cycles by applying Fourier transforms. The Fourier transforms yield amplitude representations of the angle trajectories. We consider the frequencies corresponding to a multiple of the number of gait cycles as regular and the other frequencies as irregular, i.e., a frequency f is *regular* if and only if

$$f \mod \#\text{double-steps} = 0. \qquad (2)$$

Otherwise, f is called *irregular*. Here, we use a normalization of the frequencies with respect to the execution time of a sequence consisting of five double steps. We then set the amplitudes of the irregular frequencies to zero. Finally, we perform an inverse Fourier transform from the amplitude spectrum based representation into the data points representation.

4.2.5 Uniform Trajectory Lengths

For our optimization, it is necessary that the human's and the robot's trajectory have the same length, i.e., the same number of data points. We use the joint angle trajectories of both knees and match the temporal position of the peaks in the robot's and the human's knee motion. Then, we perform a linear interpolation to estimate data points and to achieve a uniform trajectory length. We choose the length of the robot's trajectory as the reference length.

5 GAIT OPTIMIZATION

The input to the optimization process are the robot's and the human's joint angle trajectories which are computed as described in the previous section. The goal of the optimization is to adapt the robot's joint angle trajectories towards those of the human.

In an iterative fashion, we adapt the individual joint angle trajectories over the entire walking sequence. In each iteration, the relative change of an individual joint angle is the same for all time steps of the walking sequence. For each time step, the new angle of an individual joint i is computed proportional to the difference between the robot's original joint angle θ_i^{rob} and the human's joint angle θ_i^{hum} at that time step. Using the change factor $0 \leq \alpha_i \leq 1$ for joint i, the joint's new angle $\hat{\theta}_i^{rob}$ is computed as

$$\hat{\theta}_i^{rob} \leftarrow \theta_i^{rob} + \alpha_i(\theta_i^{hum} - \theta_i^{rob}). \tag{3}$$

Here, we omitted the time index due to readability.

Our optimization methods presented in the following aim at adapting α_i for each joint i so that a stable walk is achieved that is as similar as possible to the human's gait. Our criterion how to assess stability based on data of the pressure sensors in the robot's feet is detailed in Sec. 6.1. As similarity measure, we use a quantity which is based on the distance between the human's and the robot's joint angle trajectories (the distance is computed using Eq. 1). In particular, we compute in each iteration of the optimization the similarity between the robot's current gait $\hat{\boldsymbol{M}}^{\boldsymbol{rob}}$ and the human's gait $\boldsymbol{M}^{\boldsymbol{hum}}$ given the robot's initial gait $\boldsymbol{M}^{\boldsymbol{rob}}$ as

$$\text{similarity}(\hat{\boldsymbol{M}}^{\boldsymbol{rob}}, \boldsymbol{M}^{\boldsymbol{hum}}) = \frac{\text{dist}(\boldsymbol{M}^{\boldsymbol{rob}}, \hat{\boldsymbol{M}}^{\boldsymbol{rob}})}{\text{dist}(\boldsymbol{M}^{\boldsymbol{rob}}, \boldsymbol{M}^{\boldsymbol{hum}})}. \tag{4}$$

A similarity value of 0 means that there was no progress towards the human gait, while a value 1 stands for a perfect match. We consider the optimization to have converged when the gain in the similarity is below a certain threshold.

In the following, we present the optimization methods we investigated. According to the DOFs of the *HOAP-2*, the search space of the optimization is 11-dimensional since symmetric joints are always changed equally and the hand joints as well as the two head joints are not considered as relevant.

5.1 Single Component Sampling

First, we considered the simple approach of sampling one joint i in each iteration whose change factor α_i is increased according to the change rate $\Delta\alpha_i > 0$:

$$\alpha_i \leftarrow \alpha_i + \Delta\alpha_i \tag{5}$$

When the gait resulting from the new change factor leads to a stable gait, we keep the increased change factor. If not, we refuse the change. If a certain number of attempts to increase the change factor of joint i fail to generate a stable gait, we halve the change rate $\Delta\alpha_i$. Initially, the change rate is the same for all joints. Note that in this approach, all resulting joint angle trajectories are at least as similar to the human as the original trajectories.

5.2 Correlation-based Optimization

The next approach considers the correlation between the trajectories of the individual joints and allows for changing all joints simultaneously. The idea behind this technique is to take into account that the movements of certain joints are related to each other. For example, the movement of a leg during walking corresponds to several joint angle changes, such as the ankle, knee, and hip joint.

During the optimization process, we therefore change the trajectory of each joint according to the correlation between this joint and the sampled joint. The correlation between the robot's joints $\boldsymbol{Cor}^{rob}$ and the human's joints $\boldsymbol{Cor}^{hum}$ are learned for the robot from the initial gait and for the human from the recorded trajectories. To compute these correlations between joints, we consider the individual joint angle trajectories as data points. The symmetric matrix $\boldsymbol{Cor}$ then contains the correlation between any two of the joint angle trajectories.

The pseudocode of the correlation-based optimization can be found in Alg. 1. Since the correlation of the robot's and the human's joints can be different, we perform a linear interpolation between these two values according to the current change factor. *generateJointAngleTrj* then computes the new joint angle trajectories for the robot according to Eq. 3. As before, we start with the same change

Algorithm 1 Correlation-based Optimization.
$correlationBasedOptimization(\boldsymbol{M}^{rob}, \boldsymbol{M}^{hum}, \Delta\boldsymbol{\alpha})$

Input: Joint angle trajectories of the robot's gait $\boldsymbol{M}^{rob}$ and of the human's gait $\boldsymbol{M}^{hum}$, change rate for the N individual joint angles $\Delta\boldsymbol{\alpha} = \langle \Delta\alpha_1, \cdots, \Delta\alpha_N \rangle$.
Output: Joint angle trajectories $\boldsymbol{M}^{rob}_{opt}$ of the most human-like, stable gait found.

```
M_opt^rob ← M^rob; α ← 0;
Cor^hum ← computeCorrelation(M^hum)
Cor^rob ← computeCorrelation(M^rob)
while not converged do
    r ← random(1, ..., N)
    for all i ← 1 ... N do
        c ← (1 − α_i) · |Cor^rob_{r,i}| + α_i · |Cor^hum_{r,i}|
        α̂_i ← α_i + Δα_i · c
    end for
    M̂^rob ← generateJointAngleTrj(M^rob, M^hum, α̂)
    if evaluate(M̂^rob) = stable then
        M_opt^rob ← M̂^rob; α ← α̂
    end if
end while
return M_opt^rob.
```

rate $\Delta\alpha_i$ for all joints, but decrease this value when there is no success in several attempts in which the same joint was sampled (this is omitted in the code to ensure readability).

5.3 Hill Climbing

We furthermore developed an optimization technique based on hill climbing. This approach tries in each iteration D different change vectors $\boldsymbol{\alpha}$, which are generated by *createCandidate* and each of which contains change factors for each individual joint. The change vectors are sorted according to the associated similarity by *insertWithPriority* (which can easily be done without actually generating the joint trajectories).

We then take the change vector yielding the best stable result as starting point for the next iteration. In contrast to the previous methods, we now allow three different change rates. For each joint i, the change factor α_i is modified randomly by either $+\Delta\alpha$, 0, or $-\Delta\alpha$. In this way, there is a higher flexibility in finding configurations since changes in the opposite direction of the human are possible. The similarity between the trajectories can be computed before we evaluate the resulting gait in terms of stability. Therefore, we sort the resulting gaits according to their similarity and take the first gait evaluated as stable as starting point for the next iteration. A precondition is of course, that this gait has a higher similarity than the starting gait of this iteration. Otherwise or when no trajectories leading

Algorithm 2 Hill climbing.
$hillClimbing(\boldsymbol{M}^{rob}, \boldsymbol{M}^{hum}, \Delta\alpha, D)$

Input: Joint angle trajectories of the robot's gait $\boldsymbol{M}^{rob}$ and of the human's gait $\boldsymbol{M}^{hum}$, change rate $\Delta\alpha$, D number of candidates evaluated in each iteration.
Output: Joint angle trajectories $\boldsymbol{M}^{rob}_{opt}$ of the most human-like, stable gait found.

```
M_opt^rob ← M^rob; α ← 0; P ← ∅; prev_sim ← 0
while not converged do
    for all j ← 1 ... D do
        insertWithPriority(P, createCandidate(α, Δα))
    end for
    while P ≠ ∅ do
        α̂ ← dequeue(P)
        M̂^rob ← generateJointAngleTrj(M^rob, M^hum, α̂)
        if similarity(M̂^rob, M^hum) > prev_sim then
            if evaluate(M̂^rob) = stable then
                M_opt^rob ← M̂^rob; α ← α̂; P ← ∅
                prev_sim ← similarity(M̂^rob, M^hum)
            end if
        else
            P ← ∅
        end if
    end while
end while
return M_opt^rob.
```

to a stable gait are among the test trajectories, we sample new change factors for the individual joints. If a certain number of attempts is not successful, we decrease the change rate. Alg. 2 depicts pseudocode of the hill climbing method. Initially, we use the same absolute value of the change rate $\Delta\alpha$ for all joints, which is decreased as described in Sec. 5.1 (this is omitted in the code).

5.4 Policy Gradient Optimization

Finally, we developed a method based on policy gradient estimation, which is a modification of standard policy gradient reinforcement learning [22]. The idea is to evaluate a number of change vectors $\boldsymbol{\alpha}$ in order to approximate the gradient of the similarity for each dimension of the search space and explore in the direction of the maximum. Since stability is a precondition in our application, we also take this into account during the computation of the gradient.

As in our hill climbing approach, we sample D different change vectors $\boldsymbol{\alpha}$ in each iteration to generate test gaits around the current best gait. Again, the change rate $\Delta\alpha_i$ is set randomly to $+\Delta\alpha$, 0, or $-\Delta\alpha$ for each joint i. In each iteration, all D resulting joint angle trajectories are executed and the result is added for each dimension i to one of the categories S_i^+, S_i^0, or S_i^- depending on the value of the sampled change rate for this joint. To be more specific, the progress with respect to the similarity (see Eq. 4) is added for joint i to S_i^+ if $\Delta\alpha_i$ is positive, otherwise it is added to S_i^- or S_i^0, respectively. If the resulting gait is instable, 0 is added to the respective category. After evaluating each of the test gaits, the average scores Avg_i^+, Avg_i^0, and Avg_i^- of S_i^+, S_i^0, and S_i^- are computed. The corresponding values give an estimate of the gain in changing the joint angle θ_i in this specific direction and indicate how θ_i should be changed to improve the result. For each joint i, a change factor δ_i is computed in each iteration as follows:

$$\delta_i \leftarrow \begin{cases} 0 & \text{if } Avg_i^0 > Avg_i^+ \text{ and} \\ & \quad Avg_i^0 > Avg_i^- \\ Avg_i^+ - Avg_i^- & \text{otherwise} \end{cases} \tag{6}$$

Afterwards, δ_i is normalized and multiplied by the step size η, which is a scalar. The final change factor α_i for joint i is then determined as:

$$\alpha_i \leftarrow \alpha_i + \eta \cdot \frac{\delta_i}{|\boldsymbol{\delta}|} \tag{7}$$

Thus, each joint is changed towards the direction yielding the highest improvement with η determining the amount of change. Also in this optimization approach, the change rate $\Delta\alpha$ is decreased in case of no success.

6 EXPERIMENTS

We performed extensive experiments to evaluate the different optimization methods. Note that all methods will converge to a local optimum, which is not necessarily the global one. We carried out ten different experiments for each of the learning methods in order to perform a significance analysis, i.e., we performed ten restarts for each optimization technique.

To evaluate the stability of the resulting gait in each individual iteration of the optimization, the robot executed five gait cycles with identical joint angle trajectories. In the beginning of each of these runs, the execution started with the same stable motions (the initial ones, which were obtained via the CPG) that were smoothly transformed into the desired movements corresponding to the joint angle trajectories from the optimization process. Thus, we modified the initial gait into the desired one by adapting the joint angle trajectories moderately over a certain time interval. In this way, we avoided instability resulting from large changes in the beginning. Then, the robot executed five gait cycles with identical joint angle trajectories. At the end of each run, the trajectories were smoothly adapted to stand still stably (similar to the beginning of the run). In case the robot fell, the run was aborted.

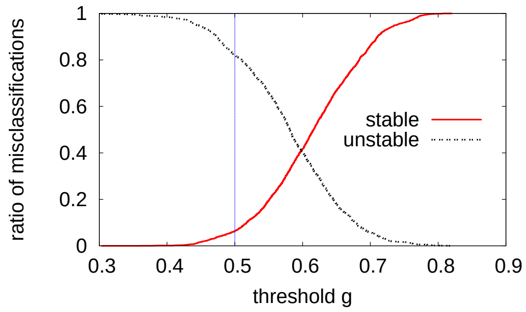

Fig. 4. Relative number of misclassifications according to a given threshold g. The red (solid) line shows the amount of stable motions which are classified as instable according to g. The black (dashed) line shows the amount of instable runs, which are classified as stable. We chose a value of 0.5 for g.

6.1 Stability Evaluation

Since the execution of the trajectories in simulation is noisy and varies in the different runs, it cannot be guaranteed that a gait is stable if the robot does not fall in a single run. For example, the robot's feet may not always rest precisely on the ground. To assess whether executed joint angle trajectories result in stable motions of the robot, we therefore use two criteria. First, we carry out five repetitions of the same executed motions. We consider a gait as instable, as soon as the robot falls in a run.

To further analyze the gaits that were not classified as instable so far, we use a second criterion. This is based on the assumption that a stable motion should be regular in the sense of foot contacts on the ground. To measure the degree of regularity of a run, we use the pressure sensors in the robot's feet and process it in the following way: First, we transform the data of each sensor into the amplitude spectrum using the Fourier transform. Then, we compute the ratio of the sum of the amplitudes of the regular (as defined in Eq. 2) versus the sum of the irregular frequencies. Since we use two sensors in each foot, we get a total of four ratios. We use the worst, i.e., the smallest value to classify the regularity of the run. If the value drops below a given threshold g, we consider the run as instable. Obviously, using this criterion, stable motions can possibly be classified as instable. However, it also decreases the number of instable results that are falsely classified as stable. Fig. 4 shows the classification result on a test set depending on g. While a low threshold filters out fewer instable motions, it misclassifies only few stable gaits. For our experiments, we chose a value of 0.5 for g,which we found out to yield good results.

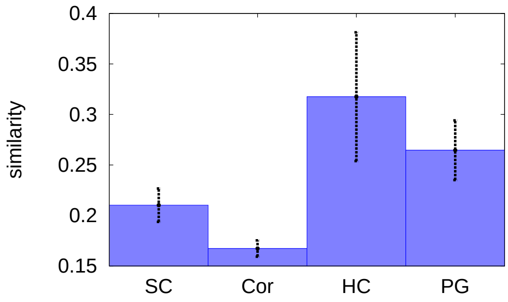

Fig. 5. Comparison of the different optimization methods. Hill climbing (HC) and policy gradient estimation (PG) perform significantly better than the single component (SC) and correlation-based (Cor) methods. All methods are able to significantly improve the similarity to the human gait. Shown are the mean and the 95% confidence interval. Note that a similarity of 0 corresponds to the robot's initial gait.

6.2 Comparison of the Optimization Methods

The parameters used for the presented experiments are the following: The single component method used $\Delta\alpha = 0.05$, the correlation based method $\Delta\alpha = 0.1$, the hill climbing method $\Delta\alpha = 0.1$ and $D = 10$, and the policy gradient approach $\Delta\alpha = 0.1$, $D = 25$, and $\eta = 0.2$. Note that the $\Delta\alpha$ values may decrease in case no improvement is achieved by the optimizations after five attempts.

Fig. 5 shows the results obtained with the different optimization techniques. As can be seen, all methods are able to significantly improve the initial robot's gait. Hill climbing and policy gradient optimization significantly outperform the two other methods. In comparison to single component sampling, the advantage of these methods is that several joints are adapted simultaneously. In policy gradient optimization, the gradient can only be approximated inaccurately and, therefore, it does not perform better than hill climbing.

Fig. 6 depicts the evolution of the change factor for different joints over time. The results are shown for the most human-like, stable gait found with hill climbing. As can be seen, especially the knee joint can be highly adapted. With hill climbing, the lateral swing can be reduced up to 12.23%. In Fig. 7, snapshots of the initial (left image) and the optimized gait (center image) are depicted.

Interestingly, in several experiments with the hill climbing and the policy gradient estimation, it came out that the robot is leaned forward (see right image of Fig. 7). This is similar to the behavior of humans when carrying heavy loads. Thus, the robot learned how to walk stably and human-like despite the fact that it does not explicitly know about the different weight distribution.

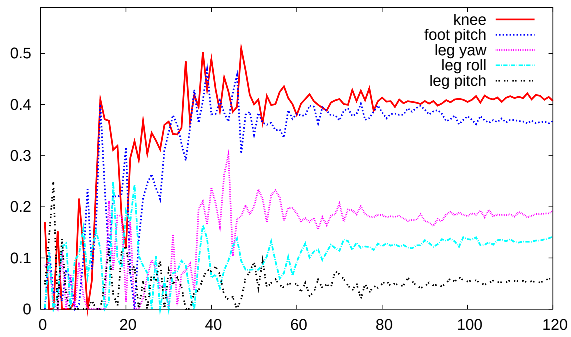

Fig. 6. Evolution of the change factor of selected joints over time for the best gait found with hill climbing.

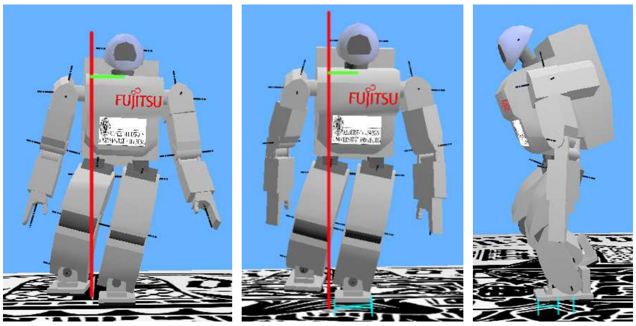

Fig. 7. Comparison of the lateral swing of the initial (left) and the optimized gait (center). The robot learned to lean forward such as humans do when carrying heavy loads (right).

7 CONCLUSIONS AND FUTURE WORK

In this article, we presented an approach to generate human-like walking patterns for humanoid robots. We investigated four optimization techniques, which work on joint angle basis and do not need any further knowledge. Our optimization starts with a gait for the humanoid obtained by a central pattern generator and tries to optimize it with respect to human resemblance in terms of joint angle difference. For the optimization, we use human data recorded by a full-body motion capture system.

We demonstrated in simulation experiments with a *HOAP-2* robot that all presented optimization techniques are able to generate joint angle trajectories that are significantly more human-like than the robot's original gait. The best performance showed methods based on hill climbing and policy gradient estimation. We achieved these good results despite different anatomy and weight distribution of the robot and the human and without explicitly modeling these aspects.

A drawback of this approach, which is solely based on joint angle trajectories, is that the trajectories cannot be modified in a temporal manner. For example, this could be necessary when the motions of the arms yield an acceleration that has an impact on the whole body dynamics and could be used to compensate for other influences. A possible solution could be to take additional means in the optimization into account that allow for shifting or morphing the trajectories over time. In future work, we will also analyze the effects of adapting the frequency of the robot's gait since this highly influence the stability.

Furthermore, it should be investigated whether the resulting robot behavior leads to increasing social acceptance. In our research, we concentrated on a limited set of aspects that have an effect on the perceived human-likeness. Further investigations should be concerned with secondary properties such as the position of body segments or a jerk-free motion. Additional psychological tests could try to find the most important factors to receive a high social acceptance.

ACKNOWLEDGMENT

This work has partly been supported by the German Research Foundation (DFG) under contract number SFB/TR-8. We would like to thank Christian Liedtke and Slawomir Grzonka for their effort in collecting data with the MVN motion capture system.

REFERENCES

[1] G. Endo, J. Morimoto, J. Nakanishi, and G. Cheng, "An empirical exploration of a neural oscillator for biped locomotion control," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2004.

[2] S. Aoi and K. Tsuchiya, "Locomotion control of a biped robot using nonlinear oscillators," *Journal of Autonomous Robots*, vol. 19, no. 3, pp. 219–232, 2005.

[3] M. Vukobratovic and B. Borovac, "Zero-moment point – Thirty five years of its life," *Int. Journal of Humanoid Robotics*, vol. 1, no. 1, pp. 157–173, 2004.

[4] F. Faber and S. Behnke, "Stochastic optimization of bipedal walking using gyro feedback and phase resetting," in *Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2007.

[5] C. Niehaus, T. Röfer, and T. Laue, "Gait-optimization on a humanoid robot using particle swarm optimization," in *2nd Workshop on Humanoid Soccer Robots at the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2007.

[6] T. Hemker, M. Stelzer, O. von Stryk, and H. Sakamoto, "Efficient walking speed optimization of a humanoid robot," *Int. Journal of Robotics Research*, vol. 28, no. 2, pp. 303–314, 2009.

[7] Q. Huang, Z. Yu, W. Zhang, X. Duan, Y. Huang, and K. Li, "Generation of humanoid walking pattern based on human walking measurement," in *Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2008.

[8] R. Chalodhorn, D. Grimes, K. Grochow, and R. Rao, "Learning to walk through imitation," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2007.

[9] H. Serhan, C. Nasr, Hénaff, P., and F. Ouezdou, "A new control strategy for ROBIAN biped robot inspired from human walking," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008.

[10] Cyberbotics, "Webots 6." http://www.cyberbotics.com/, 2009.

[11] L. Fujitsu Automation Co., "HOAP-2 instruction manual." http://robita.iiita.ac.in/hoap2instruction03e.pdf, 2004.

[12] T. Geng, B. Porr, and F. Wörgötter, "Fast biped walking with a sensor-driven neuronal controller and real-time online learning," *Int. Journal of Robotics Research*, vol. 25, no. 3, pp. 243–259, 2006.

[13] D. Lizotte, T. Wang, M. Bowling, and D. Schuurmans, "Automatic gait optimization with gaussian process regression," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2007.

[14] N. Kohl and P. Stone, "Machine learning for fast quadrupedal locomotion," in *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2004.

[15] J. Denk and G. Schmidt, "Walking primitive databases for perception-based guidance control of biped robots," *European Journal of Control*, vol. 13, no. 2-3, pp. 171–188, 2007.

[16] K. Harada, K. Hauser, T. Bretl, and J.-C. Latombe, "Natural motion generation for humanoid robots," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2006.

[17] F. Faber, M. Bennewitz, and S. Behnke, "Controlling the gaze direction of a humanoid robot with redundant joints," in *Proc. of the 17th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2008.

[18] J. E. Bobrow, B. Martin, G. Sohl, E. C. Wang, J. Kim, and F. C. Park, "Optimal robot motions for physical criteria," *Journal of Robotic Systems*, vol. 18, no. 12, pp. 785–795, 2001.

[19] M. Svinin, Y. Masui, Z. Luo, and S. Hosoe, "On the dynamic version of the minimum hand jerk criterion," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[20] XSENS Technologies B.V., "MVN: Inertial Motion Capture." http://xsens.com/en/entertainment/mvn, 2009.

[21] J. Kruskall and M. Liberman, "The symmetric time warping problem: From continuous to discrete," in *Time Warps, String Edits and Macromolecules*, pp. 125–161, Reading, USA: Addison-Wesley Publishing Co., 1983.

[22] R. Sutton, D. Mcallester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," *Advances in Neural Information Processing Systems*, vol. 12, pp. 1057–1063, 2000.

Sven Wehner studied Computer Science at the University of Freiburg. In 2009 he got his B.Sc. and is now heading for the Master's degree. His research focuses on navigation techniques for humanoids as well as for wheeled robots.

Maren Bennewitz studied Computer Science at the University of Bonn and received her Diploma degree in 1999. From 1999-2004, she was a Ph.D. student at Freiburg University. In 2004 she joined the humanoid robots lab which she heads since October 2008 as an assistant professor. Her research interests include humanoid robots, state estimation, robust navigation, and human-robot interaction.

AUTHORS' ADDRESSES

Sven Wehner, B.Sc.
Prof. Dr. Maren Bennewitz,
Humanoid Robots Lab,
University of Freiburg,
Georges-Koehler-Allee 79, 79110 Freiburg, Germany
email: wehner@informatik.uni-freiburg.de,
maren@informatik.uni-freiburg.de

Received: 2010-03-19
Accepted: 2011-02-02