Model Predictive Control with a Visuomotor System for Physics-based Character Animation

HAEGWANG EOM, Visual Media Lab, KAIST and Weta Digital
DASEONG HAN, Handong Global University
JOSEPH S. SHIN (formerly SUNG YONG SHIN), Handong Global University and KAIST
JUNYONG NOH, Visual Media Lab, KAIST

This article presents a Model Predictive Control framework with a visuomotor system that synthesizes eye and head movements coupled with physics-based full-body motions while placing visual attention on objects of importance in the environment. As the engine of this framework, we propose a visuomotor system based on human visual perception and full-body dynamics with contacts. Relying on partial observations with uncertainty from a simulated visual sensor, an optimal control problem for this system leads to a Partially Observable Markov Decision Process, which is difficult to deal with. We approximate it as a deterministic belief Markov Decision Process for effective control. To obtain a solution for the problem efficiently, we adopt differential dynamic programming, which is a powerful scheme to find a locally optimal control policy for nonlinear system dynamics. Guided by a reference skeletal motion without any *a priori* gaze information, our system produces realistic eye and head movements together with full-body motions for various tasks such as catching a thrown ball, walking on stepping stones, balancing after being pushed, and avoiding moving obstacles.

CCS Concepts: • **Computing methodologies** → **Animation**; **Physical simulation**;

Additional Key Words and Phrases: 3D character animation, visuomotor system, motion control, physics-based simulation

ACM Reference format:
Haegwang Eom, Daseong Han, Joseph S. Shin (formerly Sung Yong Shin), and Junyong Noh. 2019. Model Predictive Control with a Visuomotor System for Physics-based Character Animation. *ACM Trans. Graph.* 39, 1, Article 3 (October 2019), 11 pages.
https://doi.org/10.1145/3360905

1 INTRODUCTION

Physics-based simulation has been widely used in character animation. The approach enhances physical realism in character motion and facilitates convincing responses to environmental changes. Recent studies on this topic have dealt with physically valid interactions between a character and external objects in an online manner [Hämäläinen et al. 2014, 2015; Han et al. 2014; Jain et al. 2009]. Most of these studies assumed that the character can access all information necessary to interact with objects, such as their precise trajectories. This often leads to unrealistic behaviors that humans would not exhibit in real life. For example, a character would avoid an unseen ball flying to him/her from behind, which cannot be achieved by humans in real life. This calls for a new framework of motion synthesis that couples physics-based simulation with visual perception.

In general, a visuomotor coordination entails essential secondary behaviors such as head and eye movements to adjust visual attention. For instance, an observer pays attention to nearby obstacles while walking, to avoid collision by turning the head and eyes toward these. Despite the significance of these human behaviors, few papers have addressed physics-based motion synthesis for a full-body character equipped with a vision system, partly due to difficulties in modeling visual perception of humans as well as in achieving robust real-time physics-based character control. The control problem with an integrated visuomotor system leads to a Partially Observable Markov Decision Process (POMDP) [Sondik 1971], since the character's action results from estimated environment states under certain uncertainty, just like a human's action performed based on the information with uncertainty, which is instantaneously gathered through the eyes. Several studies have shown that behaviors similar to those of humans could be modeled with POMDPs [Baker et al. 2017; Belousov et al. 2016], one particular example being eye–hand coordination [Erez and Smart 2009]. In general, however, optimal control on the POMDP is very difficult except for simple cases [Madani et al. 2003; Papadimitriou and Tsitsiklis 1987].

To address this issue, we present a novel Model Predictive Control– (MPC) based framework with a visuomotor system that integrates visual perception and full-body motion control. This integration allows the synthesis of full-body motions of a character that can naturally interact with external objects based on a partial information on the environment. For subtle coordination between a vision system and the full-body dynamics, it is necessary to synthesize a physically correct full-body motion in an on-line manner and also to explore the environment through the vision system that drives eye/head movements. Note that this coordination is extremely difficult, if not impossible, to achieve with a kinematic-based vision system.

We adopt a simplified POMDP to control our visuomotor system effectively. Inspired by the work of Erez et al. [2012], we model our

Haegwang Eom work conducted while at Visual Media Lab, KAIST.
This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2017R1A2A1A05000979).
Authors' addresses: H. Eom, Visual Media Lab, KAIST and Weta Digital; email: gongguri858@gmail.com; D. Han, Handong Global University; email: dshan@handong.edu; J. S. Shin (formerly Sung Yong Shin), Handong Global University and KAIST; email: josephshin@gmail.com; J. Noh (corresponding author), Visual Media Lab, KAIST; email: junyongnoh@kaist.ac.kr.

0730-0301/2019/10-ART3 $15.00
https://doi.org/10.1145/3360905

POMDP as a fully observable belief MDP using a Kalman Filter and further make its belief update deterministic, which simplifies the formulation of trajectory optimization for the visuomotor system. The role of the Kalman filter is to mimic the behavior of humans [Rao and Ballard 1997; Welch and Bishop 2006] in tracking objects under uncertainty. Unlike the previous approach, we deal with a much more complex control problem based on full-body dynamics with contacts and a vision system that reflects the characteristics of human eyes more faithfully, exploiting focal length adjustment, saccades and pursuits.

The contributions of this article are summarized as follows: We propose a novel MPC-based framework with a visuomotor system that effectively integrates human vision and full-body dynamics models. Formulating the control problem of our visuomotor system with a POMDP, we simplify it with a deterministic MDP. Guided by reference motion data without gaze behaviors such as head and eye movements, our framework can produce full-body locomotive motions with natural looking secondary behaviors such as head turning, eye saccades, and pursuits. Our framework can deal with multiple objects by automatically switching the point of sight through the optimization that reduces the uncertainty of our vision system on the moving objects approaching to the character.

2 RELATED WORK

Physics-based character control for realistic motion synthesis is a challenging issue in computer graphics due to its complexity and high dimensionality. Early studies in physics-based full-body motion synthesis employ constrained trajectory optimization [Fang and Pollard 2003; Liu et al. 2005; Mordatch et al. 2012; Popović and Witkin 1999; Witkin and Kass 1988]. Ye and Liu [2010] and Han et al. [2014] adopt Differential Dynamic Programming (DDP) to obtain a locally optimal solution for a motion control problem while achieving an interactive performance. Tessa et al. [2012] approximate the contact dynamics of a full-body robot with smoothing functions so that the resulting dynamics equations become differentiable and solvable analytically. Todorov [2014] presents an improved version of the smoothed contact dynamics based on the regularization of contact impulses. Recently, Han et al. [2016] propose an interactive data-guided MPC-based framework for full-body character control, accelerated with efficient techniques for computing derivatives of the system dynamics. These approaches have not taken into account a vision system and resulting gazing behaviors. Built on a full-body MPC framework as in Han et al. [2016] our system focuses on generating secondary behaviors for gazing, including eye and head motions, to synthesize convincing and more human-like full-body motions.

To facilitate realistic interaction of a character with an environment, Jain et al. [2009] present a framework that can automatically generate responsive motions to environmental changes. They employ a visual sensor model to measure the proximity of environment objects from a character. Han et al. [2014] apply a low-dimensional model to their MPC-based framework for legged locomotion synthesis and demonstrate its effectiveness on avoiding obstacles and walking on stepping stones. Hämäläinen et al. [2014, 2015] propose sampling-based MPC schemes for balance recovery and interaction with the environment, with improved computational efficiency compared with previous sampling-based methods [Al Borno et al. 2013; Liu et al. 2015, 2010]. These studies assume that the character knows the environmental information fully from the beginning, unlike our system in which the character gathers the environmental information with uncertainty through visual observation.

During the Past decade, studies on gaze animation of virtual agents have been conducted actively in computer animation and virtual reality for realistic interaction of the agents with the environment. Ruhland et al. [2015] provide an excellent survey on synthesizing gaze behaviors in computer animation. Early work on gaze animation synthesis [Deng et al. 2005; Lee et al. 2002] focuses on saccades and pursuits [Leigh and Zee 2015]. Lance et al. [2010] suggest a method to synthesize gaze behaviors, including eyes, head, and upper body movements. Neog et al. [2016] demonstrate real-time animation of realistic soft tissue movement around eyes in accordance with input gaze behaviors. Peters et al. [2010] and Pejsa et al. [2015] propose a procedural model for the control of eye movement based on neurophysiological observations. Recently, Pejsa et al. [2016] present a method for adding gaze animation to existing skeletal full-body motion data, considering scene features and the head and torso kinematics. Such studies mainly layer gaze behaviors on given motions in an off-line manner. Unlike these studies, our approach synthesizes gaze behaviors in accordance with full-body dynamics while responsively reacting to the time-varying environment in an on-line manner.

A POMDP-based approach [Thrun et al. 2005] has been utilized to model a decision process for planning under uncertainty, such as robot control [Pineau and Gordon 2007] and a baseball outfielder's ball-tracking [Belousov et al. 2016]. However, this approach has been mainly applied to problems in a low-dimensional and discrete space because of its computational burden [Papadimitriou and Tsitsiklis 1987]. To address this issue, approximation schemes for the belief dynamics have been proposed [Du Toit and Burdick 2012; Platt Jr et al. 2010; Van Den Berg et al. 2012, 2017]. Kalman filters have been widely used for state estimation of non-linear systems [Welch and Bishop 2006]. Rao and Ballard [1997] employ a Kalman filter for visual cortex modeling. Building the vision system on top of a Kalman filter, we adopt a similar approximation scheme to incorporate this system into the physics-based full-body motion control.

In character animation, there have been a few studies on motion synthesis based on the internal models of a human in perceiving external objects. Yeo et al. [2012] propose a kinematic upper-body motion synthesis system that is capable of perceiving an object using a vision system. Recently, Nakada et al. [2018] adopt a vision deep neural network for the biomimetic sensorimotor control of muscle dynamics. Unlike these earlier studies that handle only a single object, our system can visually track multiple static or moving objects by automatically selecting gazing targets and switching between these based on the uncertainty model that we adopt for our vision system. In addition, our system can generate various types of full-body motions such as running or balancing, in which the root of a character moves dynamically. Erez et al. [2012, 2011] present an optimal control approach to eye-hand coordination, based on a minimax algorithm and an MDP with

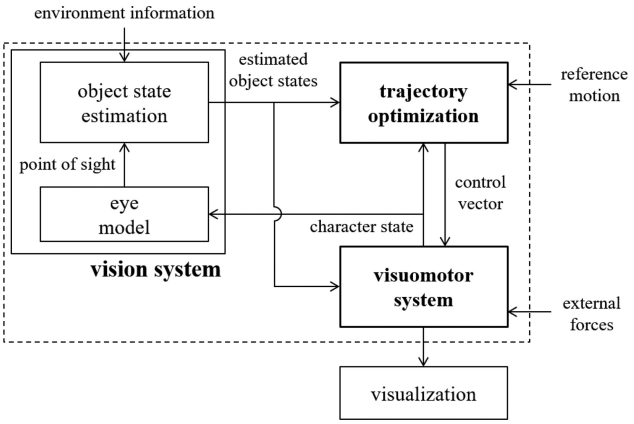

Fig. 1. System overview.

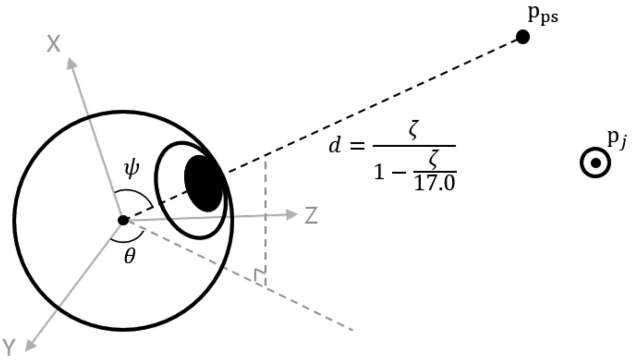

Fig. 2. Eye coordinate frame.

belief approximation. They demonstrate the effectiveness of their approach for a simplified 2D model with an eye and two hands. We extend these methods to full-body character control based on articulated-body dynamics with contacts.

3 SYSTEM OVERVIEW

As shown in Figure 1, our framework consists of three components: a vision system, a trajectory optimizer, and a visuomotor system. Given a full-body character state including an eye pose from the visuomotor system, the vision system determines the character's point of sight and estimates the states of perceived objects by approximating the belief dynamics of the POMDP using a Kalman filter to track the noisy perceived states of external objects. The estimated state of each object is composed of its position and velocity together with their uncertainty, which is used by the trajectory optimizer to generate a locally optimal control policy that adjusts full-body motion, eye and head movements, and focal length, guided by the reference motions. We adopt DDP [Jacobson and Mayne 1970] to find the control policy of a full-body character. To support an online interactive performance, our system repeatedly updates the character control policy for a short time window while shifting it along the time axis. The visuomotor system performs a forward dynamics simulation to update the full-body character state, given the control policy together with external forces, if there is any.

4 VISION SYSTEM

This section describes our vision system. In Section 4.1, we model an eye coordinate frame and eye behaviors such as saccades and pursuits. In Section 4.2, we describe how to determine the point of sight at which the character is looking. In Section 4.3, we discuss how to estimate the state of a moving object with our vision system.

4.1 Eye Model

As shown in Figure 2, we adopt a spherical coordinate system to define the eye frame. In this frame, we use three parameters to specify an eye pose $\mathbf{e} = [\theta\, \psi\, \zeta]^{\mathrm{T}}$. Here θ and ψ are respectively the azimuthal and polar angles, which together represent a gaze direction, and ζ is the focal length. We bound each parameter with a normal range of eye movement that can be obtained from medical research [Serway et al. 2018; Shin et al. 2016]: That is, $-44.2^\circ \le \theta \le 44.2^\circ$ for adduction and abduction, $-47.1^\circ \le \psi \le 27.7^\circ$ for depression and elevation, and $0.0 \le \zeta \le 17.0$ (mm) for focal length. The focal length approaches 17.0mm when human looks at a point at infinity.

We also bound eye movement to generate realistic gaze behaviors, saccades and pursuits, which are important characteristics of human eyes. Saccades and pursuits refer to rapid movements of eyeballs to find new objects and slow eyeball movements to track objects, respectively. As in Yeo et al. [2012], we adopt a simplified profile of saccades and pursuits and impose velocity constraints on these behaviors based on the results from Robinson et al. [1965], Meyer et al. [1985], Leigh and Zee [2015], and Itti et al. [2006]. Note that our system only uses velocity constraints for a natural gaze behavior and does not enforce explicit saccades or pursuits in relation to the object motion.

In particular, we set the maximum saccade speed to 800°/s and the maximum pursuit speed to 100°/s considering only azimuthal and polar eyeball movements. Saccades repeat in every 200ms interval, each followed by a 200ms pause for recharge. Pursuits occur simultaneously even during the recharge times for saccades. We incorporate all of these in our vision system as follows:

$$\left\|\dot{\mathbf{e}}_{\theta,\psi}\right\| \le \mathcal{B},$$

where $\dot{\mathbf{e}}_{\theta,\psi}$ is the eyeball speed for azimuthal and polar movements and $\mathcal{B}$ is the upper bound of its magnitude.

4.2 Point of Sight

The point of sight $\mathbf{p}_{\mathrm{ps}}$ plays a key role in integrating the full-body dynamics with the vision system, because $\mathbf{p}_{\mathrm{ps}}$ is a function of the full-body joint position $\mathbf{q}$ and eye pose $\mathbf{e}$:

$$\mathbf{p}_{\mathrm{ps}}(\mathbf{q}, \mathbf{e}) = \mathbf{p}_{\mathrm{head}}(\mathbf{q}) + \mathrm{W}_{\mathrm{head}}(\mathbf{q})\, \mathbf{p}_{\mathrm{eye}} + d(\mathbf{e})\hat{\mathbf{z}}(\mathbf{q}, \mathbf{e}).$$

Here $\mathbf{p}_{\mathrm{head}}$ and $\mathrm{W}_{\mathrm{head}}$ are the position and the rotation matrix of the head with respect to the global frame, respectively. $\mathbf{p}_{\mathrm{eye}}$ is the midpoint between the two eyes with respect to the head frame, and $d(\mathbf{e}) = \zeta/(1 - \zeta/17.0)$ is the distance between $\mathbf{p}_{\mathrm{eye}}$ and the point of sight. $\hat{\mathbf{z}}$ is a unit vector pointing to the gaze direction, that is, the z -axis extracted from matrix $\mathrm{W}_{\mathrm{eye}}(\mathbf{e})\mathrm{W}_{\mathrm{head}}(\mathbf{q})$. $\mathrm{W}_{\mathrm{eye}}$ is the rotation matrix of $\mathbf{p}_{\mathrm{eye}}$ with respect to the head frame.

4.3 Object State Estimation

In this section, we discuss how to estimate the object state from a partial observation. To generate realistic gaze behaviors, we imitate the human vision system through estimating the object state under uncertainty instead of using the true full state of an object. First, we introduce an observation model that measures the object position from a true object state obtained through the physics-based simulation. Then, we discuss a process model that produces the predicted object state. Finally, we describe how to estimate the object state based on the belief update using a Kalman filter [Erez and Smart 2012].

Under the assumption that the character does not know the true state, our system makes an observation from the true state of an object. To construct an observation model that abstracts human visual sensors, we assume that humans perceive the positional information of a moving object to estimate its velocity instantaneously. Moreover, the positional information is not accurate in general, particularly when the object is distant from the point of sight. Under these assumptions, we formulate the observation model of our vision system as follows:

$$\begin{aligned} & \mathbf{z}_j = \mathbf{p}_j + \rho = \mathrm{H}\mathbf{x}_j + \rho, \\ \text{where} \quad & \mathrm{H} = [\,\mathrm{I}\,\mathrm{O}\,],\ \rho \sim \mathcal{N}(0, \mathrm{R}_j), \\ \text{and} \quad & \mathrm{R}_j = r_j * \mathrm{I}, \\ \text{where} \quad & r_j = \left(1 - e^{\frac{-\|\mathbf{p}_{\mathrm{ps}} - \mathbf{p}_j\|}{2\eta}}\right). \end{aligned} \tag{1}$$

Let N_{obj} be the number of objects in the environment that are perceived by the vision system. Our observation model makes partial observation $\mathbf{z}_j$ on the state of an object from its true state $\mathbf{x}_j = [\mathbf{p}_j{}^{\mathrm{T}}\ \mathbf{v}_j{}^{\mathrm{T}}]^{\mathrm{T}}$ that is composed of positon $\mathbf{p}_j$ and velocity $\mathbf{v}_j$, considering observation noise ρ, for $j = 0, 1, \ldots, N_{\mathrm{obj}} - 1$. H is the transformation matrix that maps the object state onto a noise-free partial observation, that is, the position of an object. Observation noise ρ has a multivariate Gaussian distribution with covariance R_j. η is the constant for the size of the fovea. Covariance matrix R_j is a diagonal matrix, where the value of diagonal elements depends on the Euclidean distance between the point of sight $\mathbf{p}_{\mathrm{ps}}$ and the object position $\mathbf{p}_j$; specifically, the error of the observed object position increases as the character's visual attention becomes farther from the object [Erez and Smart 2012].

Under the assumption that the brain has prior knowledge on the dynamics of the object [McIntyre et al. 2001], we now formulate the process model that deals with the evolution of the object state as follows:

$$\begin{aligned} & \mathbf{x}'_j = \mathrm{F}\mathbf{x}_j + \mathbf{y} + \delta, \\ \text{where} \quad & \mathrm{F} = \begin{bmatrix} \mathrm{I} & h\mathrm{I} \\ \mathrm{O} & \mathrm{I} \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} \mathrm{O} \\ h\mathbf{a} \end{bmatrix}, \\ \text{and} \quad & \delta \sim \mathcal{N}(0, \mathrm{Q}). \end{aligned} \tag{2}$$

$\mathbf{x}'_j$ is the next state of object j, F is the state transition matrix, $\mathbf{y}$ is the control vector, h is the time difference between two consecutive observations, and $\mathbf{a}$ is the gravitational acceleration. δ is the process noise with covariance $\mathrm{Q} = \alpha\mathrm{I}$, where α is a nonnegative constant. Q represents the brain's own uncertainty of the process model. A large value of α makes the uncertainty of the estimated state of an object tend to grow quickly as the point of sight becomes farther from the object. This makes the character move the point of sight back to the object whenever the uncertainty grows sufficiently large. Varying this coefficient, we can adjust the ability of our vision system; how long the character would look at an object or when it would switch its attention from one object to another in the presence of multiple objects.

The vision system estimates the state of the object, relying on a partial observation of a true object state. Thus, the optimal control problem for this model naturally leads to a POMDP [Sondik 1971], where the underlying state of an object is partially known. We convert the POMDP into a belief MDP by introducing a probability distribution over underlying object states, known as a belief state, together with a transition function from one belief state to another, which is described as a belief update. The resulting belief MDP is a variant of the standard MDP, which takes into account belief states instead of underlying object states. In general, it is intractable to perform the belief update exactly to solve the belief MDP with an infinite-dimensional belief space [Thrun 2000]. We employ a Kalman filter [Erez and Smart 2012; Sardag and Akin 2006; Van Den Berg et al. 2012] to approximate the belief update as a Gaussian update [Stengel 1994] based on a partial observation on the external objects in the environment.

We start with defining a belief state b over underlying states of an external object as a Gaussian distribution, that is, $b(\mathbf{x}_j) = \mathcal{N}(\mu_j, \Sigma_j)$, where mean μ_j and covariance Σ_j can be regarded as the estimated state and the uncertainty of estimation, respectively. Given the current belief state b and the current observation $\mathbf{z}_j$, the vision system estimates the next belief state $b'(\mathbf{x}_j)$ through belief update $n(b, \mathbf{z}_j)$ based on the Kalman filter given in Equations (1) and (2) (see Appendix A for the full derivation), that is,

$$b'(\mathbf{x}_j) = n(b, \mathbf{z}_j) = \mathcal{N}(\mu'_j, \Sigma'_j),$$

where μ'_j and Σ'_j are the new mean and covariance at the next time step:

$$\mu'_j = \mathrm{F}\mu_j + \mathbf{y} + \mathrm{K}_j\left(\mathbf{z}_j - \mathrm{H}(\mathrm{F}\mu_j + \mathbf{y})\right), \tag{3}$$

$$\Sigma'_j = \left(\mathrm{I} - \mathrm{K}_j\mathrm{H}\right)\left(\mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q}\right). \tag{4}$$

Here $\mathrm{K}_j = (\mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q})\mathrm{H}^{\mathrm{T}}(\mathrm{H}(\mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q})\mathrm{H}^{\mathrm{T}} + \mathrm{R}_j)^{-1}$ is the Kalman gain for object j, which determines how much the current observation affects the next belief state. The belief state b will be used for the trajectory optimization later in Section 6. For compact representation, we express $\Sigma_j = \left[\begin{smallmatrix} \sigma_j^{\mathrm{p}}\mathrm{I} & \sigma_j^{\mathrm{pv}}\mathrm{I} \\ \sigma_j^{\mathrm{pv}}\mathrm{I} & \sigma_j^{\mathrm{v}}\mathrm{I} \end{smallmatrix}\right]$ as 3D vector $\sigma_j = [\sigma_j^{\mathrm{p}}\ \sigma_j^{\mathrm{v}}\ \sigma_j^{\mathrm{pv}}]^{\mathrm{T}}$.

5 VISUOMOTOR SYSTEM

In this section, we present the full dynamics of our visuomotor system that governs the evolution of the character state with an estimated object state and a character control vector. We adopt the stable articulated-body dynamics proposed by Kumar et al. [2014] and Todorov [2014]:

$$\begin{aligned} & \mathrm{M}(\dot{\mathbf{q}}' - \dot{\mathbf{q}}) = h\left(\begin{bmatrix} 0 \\ \tau \end{bmatrix} - \phi - \mathrm{B}\dot{\mathbf{q}}\right) + \mathrm{J}^{\mathrm{T}}\hat{\mathbf{f}}, \\ \text{where} \quad & \mathrm{M} = \mathrm{M}_0 + \mathrm{M}_a + h\mathrm{B}. \end{aligned}$$

Here, M is the total mass matrix, $\dot{\mathbf{q}}$ and $\dot{\mathbf{q}}'$ are respectively the joint velocities at the current and next steps, h is the integration step size, τ is the actuated joint torques, and ϕ is the bias force resulting from the gravity, Coriolis forces, and external forces if there is any. $\mathrm{B} = k_d\mathrm{I}$ is the damping gain matrix with constant k_d, and J is the kinematic Jacobian matrix at a contact point on which contact impulse $\hat{\mathbf{f}}$ is exerted. M_0 is the plain mass matrix, and $\mathrm{M}_a = \rho_a\mathrm{I}$ is the armature inertia matrix with coefficient ρ_a. For detail in building M_0, we refer the readers to the composite rigid body algorithm described in Featherstone [2008].

Let $\mathbf{x}_{\text{fb}} = [\mathbf{q}^{\mathrm{T}}\ \dot{\mathbf{q}}^{\mathrm{T}}]^{\mathrm{T}}$ and $\mathbf{u}_{\text{fb}} = \tau$ be the state and control vectors for a full-body character, respectively, where $\mathbf{q}$ is the full-body joint position. Specifically, being composed of hinge joints for elbow and knee joints and ball joints for the others, our full-body model $\mathbf{x}_{\text{fb}}$ has 40 degrees of freedom, excluding the eyeball parameters, and $\mathbf{u}_{\text{fb}}$ has 34 degrees of freedom. We derive the full-body system dynamics for the character by employing the semi-implicit Euler integration as follows:

$$\mathbf{x}'_{\text{fb}} = l(\mathbf{x}_{\text{fb}}, \mathbf{u}_{\text{fb}}) = \begin{bmatrix} \mathbf{q} + h\dot{\mathbf{q}}' \\ \dot{\mathbf{q}}' \end{bmatrix},$$
$$\text{where} \quad \dot{\mathbf{q}}' = \dot{\mathbf{q}} + h\mathrm{M}^{-1}\left(\begin{bmatrix} 0 \\ \tau \end{bmatrix} - \phi - \mathrm{B}\dot{\mathbf{q}}\right) + \mathrm{M}^{-1}\mathrm{J}^{\mathrm{T}}\hat{\mathbf{f}}.$$

Here we compute contact impulse $\hat{\mathbf{f}}$ based on a state-of-the-art smoothed contact dynamics [Todorov 2014].

By combining the full-body dynamics and the belief update in Section 4.3, we obtain the system dynamics for our visuomotor system:

$$\mathbf{x}' = g(\mathbf{x}, \mathbf{u}, \mathbf{z}) = \begin{bmatrix} l(\mathbf{x}_{\text{fb}}, \mathbf{u}_{\text{fb}}) \\ \mathbf{e} + h\dot{\mathbf{e}} \\ \mu' \\ \sigma' \end{bmatrix}.$$

Here $\mathbf{x} = [\mathbf{x}_{\text{fb}}{}^{\mathrm{T}}\ \mathbf{e}^{\mathrm{T}}\ \mu^{\mathrm{T}}\ \sigma^{\mathrm{T}}]^{\mathrm{T}}$ and $\mathbf{u} = [\mathbf{u}_{\text{fb}}{}^{\mathrm{T}}\ \dot{\mathbf{e}}^{\mathrm{T}}]^{\mathrm{T}}$ are the state and control vectors, respectively. $\dot{\mathbf{e}}$ is the time derivative of the eye pose $\mathbf{e}$. For the other symbols in the system dynamics equation, we refer the readers to Equations (3) and (4). The first two rows of the system dynamics are responsible for the update of full-body and eye states while the last two rows are responsible for the update of the current belief state, which models the process of visual perception for external objects based on partial observation $\mathbf{z}_j$. In our framework, we use the mean μ_j of the current belief states as the estimation of a perceived object state.

6 TRAJECTORY OPTIMIZATION

It is intractable to solve an optimal control problem formulated with a POMDP in general [Madani et al. 2003; Papadimitriou and Tsitsiklis 1987]. Instead of trying to find the globally optimal policy for an infinite time horizon, we adopt DDP [Jacobson and Mayne 1970] to search for a locally optimal control policy for our visuomotor system using a short time window.

DDP is, however, mainly used for deterministic systems with nonlinear system dynamics. Therefore, we further simplify our visuomotor system dynamics to handle the update function for mean μ_j, which is stochastic due to noisy observation $\mathbf{z}_j$. To obtain deterministic mean update $\hat{\mu}'_j$, we approximate μ'_j by taking the expectation utilizing the fact that Equation (3) is linear in $\mathbf{z}_j$ as follows, inspired by Erez et al. [2012]:

$$\mu'_j \coloneqq \hat{\mu}'_j = \mathrm{E}\left\{\mathrm{F}\hat{\mu}_j + \mathbf{y} + \mathrm{K}_j\left(\mathbf{z}_j - \mathrm{H}\left(\mathrm{F}\hat{\mu}_j + \mathbf{y}\right)\right)\right\} = \mathrm{F}\hat{\mu}_j + \mathbf{y}.$$

Note that the approximation $\hat{\mu}'_j$ is equivalently obtained through marginalizing belief update $n(\hat{b}(\mathbf{x}_j), \mathbf{z}_j)$ over observations $\mathbf{z}_j$:

$$\int_{\mathbf{z}_j} n(\hat{b}(\mathbf{x}_j), \mathbf{z}_j)d\mathbf{z}_j = \mathcal{N}(\hat{\mu}'_j, \Sigma'_j) \text{ where } \hat{b}(\mathbf{x}_j) = \mathcal{N}(\hat{\mu}_j, \Sigma_j).$$

This approximation is applied only for the trajectory optimization. In the actual simulation of the visuomotor system with an optimal policy, the estimated state is updated using Equation (3). As observed in Equation (4), Σ'_j is independent of observations $\mathbf{z}_j$ and thus is simply updated according to the equation.

The approximation of μ'_j results in a system of deterministic dynamics for DDP:

$$\hat{\mathbf{x}}' = \hat{\mathbf{g}}(\hat{\mathbf{x}}, \mathbf{u}) = \begin{bmatrix} l(\mathbf{x}_{\text{fb}}, \mathbf{u}_{\text{fb}}) \\ \mathbf{e} + h\dot{\mathbf{e}} \\ \hat{\mu}' \\ \sigma' \end{bmatrix}.$$

Accordingly, full-body state vector $\mathbf{x}_{\text{fb}}$ is also extended to $\hat{\mathbf{x}} = [\mathbf{x}_{\text{fb}}{}^{\mathrm{T}}\ \mathbf{e}^{\mathrm{T}}\ \hat{\mu}^{\mathrm{T}}\ \sigma^{\mathrm{T}}]^{\mathrm{T}}$, where $\hat{\mu}'$ is the matrix obtained by juxtaposing all $\hat{\mu}'_j$. Note that the state estimation update is reduced to its deterministic version by taking the most-likely observation for μ (the third row), but uncertainty σ' (the fourth row) for the object state estimation remains unchanged. The resulting system dynamics is regarded as the brain's internal model for our visuomotor system that is used to predict future full-body motions, eye movements, and belief states for external objects.

Given current state $\hat{\mathbf{x}}$ and reference skeletal trajectory $\bar{\mathrm{X}} = \{\bar{\mathbf{x}}^i_{\text{fb}} | i = 0, 1, \ldots, N-1\}$, we solve the following finite-horizon optimal control problem for control policy, $(\mathbf{u}^0, \mathbf{u}^1, \ldots, \mathbf{u}^{N-2})$ over a discrete window of size N:

$$\min_{\mathbf{u}^0, \mathbf{u}^1, \ldots, \mathbf{u}^{N-2}} \sum_{i=0}^{N-2} c(\hat{\mathbf{x}}^i, \mathbf{u}^i) + c_f(\hat{\mathbf{x}}^{N-1}),$$
$$\text{subject to} \quad \hat{\mathbf{x}}^0 = \hat{\mathbf{x}},\ \hat{\mathbf{x}}^{i+1} = \hat{\mathbf{g}}(\hat{\mathbf{x}}^i, \mathbf{u}^i) \text{ for } i = 0, 1, \ldots, N-2,$$

where $c(\hat{\mathbf{x}}^i, \mathbf{u}^i)$ and $c_f(\hat{\mathbf{x}}^{N-1})$ are the cost functions at the ith time step and the last time step, respectively. The cost function at each time step is formulated as follows:

$$c(\hat{\mathbf{x}}^i, \mathbf{u}^i) = w_{trk}c_{trk} + w_{eng}c_{eng} + w_{uct}c_{uct} + w_{bnd}c_{bnd} + w_{tsk}c_{tsk}.$$

Here w_i is the weight for the cost term c_i, where $i \in \{trk, eng, uct, bnd, tsk\}$. In what follows, we discuss each of the cost terms.

Cost for tracking: c_{trk} is the cost for tracking the reference skeletal motion and upper-body orientation, which is similar to that in Han et al. [2016]:

$$c_{trk} = \left\|\bar{\mathbf{x}}^i_{\text{fb}} - \mathbf{x}^i_{\text{fb}}\right\|^2 + \left\|\bar{\mathbf{o}}^i_{\text{ub}} - \mathbf{o}^i_{\text{ub}}\right\|^2.$$

Here $\mathbf{o}^i_{\text{ub}}$ is the displacement vector from the root to the head, and $\bar{\mathbf{o}}^i_{\text{ub}}$ is the corresponding reference vector. The first term is for tracking the reference skeletal motion and the second term is for maintaining the torso up-vector as close to that of the reference motion as possible.

Cost for energy consumption: c_{eng} is the cost for preventing a character from generating overly powerful actuation:

$$c_{eng} = \|\mathbf{u}^i\|^2.$$

Note that this term is not included in the cost function for the last step $c_f(\hat{\mathbf{x}}^{N-1})$.

Cost for uncertainty: To reduce the uncertainties of the estimated object states, we penalize the sum of the magnitudes of covariance for the belief state of each perceived object as follows:

$$c_{uct} = \sum_{j=0}^{N_{\text{obj}}-1} \left\|\sigma_j'\right\|^2.$$

This term plays an important role in tracking the states of the objects and generating convincing secondary behaviors including head and eye movements. Specifically, this term optimizes the sum of the squared diagonal elements of each observation covariance matrix R_j defined in Equation (1), which leads to the minimization of Euclidean distance r_j between point of sight $\mathbf{p}_{\text{ps}}$ and each object position $\mathbf{p}_j$. Thus, $\mathbf{p}_{\text{ps}}$ tends to move toward $\mathbf{p}_j$ in an effort to minimize r_j when there is only a single object in the environment, whether the object is moving or not. In this case, the observation of an object state by the character is guaranteed to converge to its true state as it keeps looking at the object. For detailed explanation, we refer the readers to Appendix B. Although we use a single object case to keep the explanation brief, this term itself is designed to deal with multiple objects simultaneously to move the point of sight from object to object, if necessary, guided by c_{uct}.

Constraints for eye model: Spherical coordinates θ and ψ representing the eyeball position and focal length ζ have their respective upper and lower bounds as explained in Section 4.1. Similarly, the eyeball speeds for the azimuthal and polar movements have their corresponding upper bounds.

We enforce these bounds with soft constraints as shown below:

$$c_{bnd} = \sum_{x\in\{\theta,\psi,\zeta\}} \Big(smax(x^{hi} - x;\gamma) + smax(x - x^{lo};\gamma)\Big) + smax\Big(\left\|\dot{\mathbf{e}}_{\theta,\psi}\right\| - \mathcal{B};\gamma\Big).$$

Here x^{hi} and x^{lo} are the upper and lower bounds for each eyeball coordinates and $smax(x;\gamma) = (\sqrt{x^2+\gamma^2} - x)/2$ is a smoothed version of $max(x, 0)$, which returns x if $x > 0$ and zero otherwise [Tassa et al. 2012]. Coefficient γ is used to adjust how much the function should be softened around $x = 0$.

Cost for task: c_{tsk} is a task-dependent cost term. We will explain this term in detail in Section 7.

Minimizing the objective function that consists of the above-mentioned cost terms, our system produces an optimal control policy that deals with not only full-body motions but also time-varying gaze behaviors. In particular, cost term c_{tsk} guides the full-body character to do an intended task, and cost term c_{uct} induces secondary behaviors driven by the time-varying uncertainty of each object state. Thus, our system can synthesize full-body motions with gaze behaviors to do intended tasks. As a result, full-body motions synthesized by our system would be different from those explicitly directed by the task-only cost term.

7 RESULTS

We performed four experiments to demonstrate the effectiveness our system: ball catching, walking on stepping stones, balancing after a push, and moving obstacle avoidance as shown in Figure 3. After presenting common aspects of the experiments, we discuss each of the experiments in detail. For animation results, see the supplementary video.

Implementation details

In all of the experiments, once an external object is within the character's field of view (−60°∼60°: human's field of view for binocular vision [Henson et al. 2000]) the vision system keeps track of the belief state of each object in sight as long as its uncertainty is below a certain threshold. Specifically, the system performs the trajectory optimization for the most important N_{obj} objects at every time step, which are chosen according to their own importance values depending on the task to be performed. We set the value of N_{obj} from one to four, according to an observation in cognitive science, which states that at most four objects can be tracked simultaneously in general by the human vision system [Alvarez and Franconeri 2007; Scholl and Xu 2001]. For example, in the case of obstacle avoidance, the character selects four closest obstacles at each simulation time step. Even if they are in the field of view, our system excludes unselected objects from the full-body trajectory optimization although their estimated states and uncertainties are still updated.

To solve the optimization problem presented in Section 6, we employ a state-of-the-art DDP solver equipped with acceleration techniques, such as data reuse of physical quantities of an articulated body and derivative interpolation [Han et al. 2016]. Even with the effective optimization techniques, we still need to overcome the heavy computational load for interactive control due to the high-dimensionality of our visuomotor system dynamics, where the dimensionalities of $\mathbf{x}$ and $\mathbf{u}$ are $43 + 9N_{\text{obj}}$ and 37, respectively (see Section 5).

As described in Section 4.3, the update equations for belief states (Equations (3) and (4)) are in a closed form and thus analytically differentiable. The update equation for the eye pose is also expressed in a closed form. Our method accelerates the differentiation of the system dynamics by utilizing closed-form derivatives of the update equations for eye pose and a belief state using D* [Guenter 2007] while numerically differentiating full-body dynamics $\mathbf{x}'_{\text{fb}} = l(\mathbf{x}_{\text{fb}}, \mathbf{u}_{\text{fb}})$, which cannot be expressed in a closed form. The closed-form differentiation accelerates the optimal control policy generation about four times faster than numerical differentiation. This selective closed-form differentiation becomes more effective as N_{obj} increases. All experiments were performed on a desktop computer with an Intel Core i7 processor (3.5GHz, six cores) and 32GB memory, and the resulting motions were generated at a rate of 1 to 6 frames per second without code optimization.

Depending on the task, the integration step size h is set to 0.02s or 0.03s and the window size is set to a value in a range from 0.6s to 1.2s. We set the weight for each term of the cost function as follows: $w_{trk} = 0.1$, $w_{eng} = 10^{-5}$, $w_{uct} = 1.0$, and $w_{bnd} = 0.1$. w_{tsk} was set to a value in a range from 10^{-3} to 10 depending on the task to be performed. It is time-consuming to adjust the weight values for convincing results, usually taking a few hours to a few

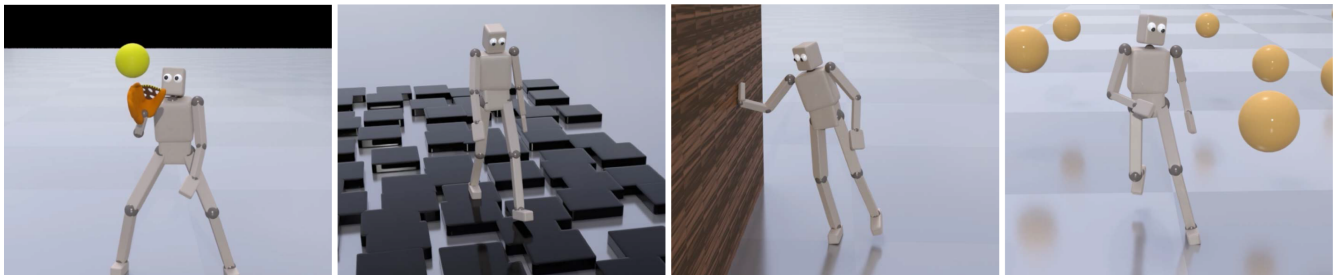

(a) Ball catching (b) Walking on stepping stones (c) Balancing by leaning on wall with one hand (d) Moving obstacle avoidance

Fig. 3. Full-body motions with eye movements produced by our MPC framework with a visuomotor system: From left to the right, catching a thrown ball, walking on stepping stones, and balancing by leaning on a wall with a hand after an unexpected push, running motion while avoiding approaching obstacles.

days. However, we believe that the proposed parameters could be a reasonable starting point when implementing similar systems.

Ball catching

We performed an experiment on our visuomotor system with a thrown ball, using a single standing pose as reference motion data. The task-dependent cost term was designed as follows:

$$c_{tsk} = \|\mathbf{p}_{\text{hand}} - \mathbf{p}_{\text{ball}}\|^2 + \left\|\mathbf{p}_{\text{ball}}^{lat} - \mathrm{p}_{\text{sp}}^{lat}\right\|^2,$$

where $\mathbf{p}_{\text{hand}}$ and $\mathbf{p}_{\text{ball}}$ are respectively the position of the catching hand and the estimated arrival position of the ball based on the current character state and the estimated ball state. $\mathbf{p}_{\text{ball}}^{lat}$ and $\mathbf{p}_{\text{sp}}^{lat}$ respectively are the lateral components of $\mathbf{p}_{\text{ball}}$ and the mid-point of the both feet approximating the center of the support polygon. The first term is responsible for moving the hand to the estimated arrival position of the ball and the second term is for moving the feet in the lateral direction if the arrival position is laterally far from the character.

As in Yeo et al. [2012], we divide the ball-catching motion into two phases: reactive and proactive. The former is the phase to take preparatory actions for catching the ball, and the latter is for taking actual actions to catch the ball. For example, a human usually lifts up the arm and keeps it lifted until the ball approaches near (the reactive phase) and then adjusts the position of the hand so as to catch the ball (the proactive phase). To produce these behaviors, our system uses different values for the weight for c_{tsk}; in particular, applying a small weight (10^{-3}) for the reactive phase and a large weight (0.1) for the proactive phase. Unlike Yeo et al. [2012], we simply divide the phase according to the estimated arrival time of the ball. The reactive phase starts at random within 1.5–2s before the arrival time, and continues until the proactive phase starts. The proactive phase begins 0.5–1s before the arrival time and lasts until the ball arrives. In our experiment, the character naturally performed reactive and proactive motions while looking at the flying ball with convincing gaze behaviors. The character also moved laterally in an automatic manner when the ball arrival position was out of reach.

We compared our framework with that proposed by Han et al. [2016], which has no vision mechanism and thus considered as a fully observable system [Spaan 2012]. Unlike in our framework, the character in the previous work had perfect knowledge on the environment and never failed in catching the flying ball, even when it is thrown from the back. To show how the uncertainty of the vision system affects gaze behaviors, we synthesized motions with different values of α (see Section 4.3): A large value of α results in a large value of Kalman gain K_j. Thus, estimated state μ_j' in Equation (3) relies more on the current noisy observation, making our system unstable and produce an unexpected result. With $\alpha = 0$, the character glanced at the flying ball once to resolve visual uncertainty and never showed more gaze behaviors. With $\alpha = 0.4$, the character failed to catch the flying ball even with steady tracking, since the uncertainty grows too rapidly. With $\alpha = 0.2$, the character exhibited realistic gaze behaviors.

To show the advantage of our approach over others in tracking multiple objects, we compared their respective produced trajectories of the point of sight. Figure 4 shows two black balls in the environment that are flying to the character. The approach that minimizes the sum of the distances from the point of sight to the balls ($\Sigma_{j=0}^{N_{\text{obj}}-1}\|\mathbf{p}_{\text{ps}} - \mathbf{p}_j\|^2$) unlike our cost term would produce similar results to ours for a single object. However, the point of sight must be placed in the middle of the two balls at each time step when the balls approach to the character, simultaneously (left), which results in few secondary behaviors. The approach of Yeo et al. [2012] would sequentially track two balls, one after the other. After tracking the first ball, it must be successful to catch it. However, the character looks at the other ball only after catching the first one, which may result in an unnatural motion because of short preparation time (middle). Our approach switches the point of sight back and forth between the two objects while reducing the total uncertainty (right). As the two balls move closer to the character, their distance becomes smaller so that the uncertainty of the ball that is farther from the point of sight does not grow rapidly any more. Thus, the point of sight stays relatively longer on the closer ball right before the character catches it. We show animation results in the supplementary video.

Walking on stepping stones

This experiment of walking on stepping stones was performed to show how gaze behaviors are produced while walking on an uneven terrain, where visual attention is important to plan foot placement accurately. To select a stone on which the stepping foot would be placed, our system computed which stone was

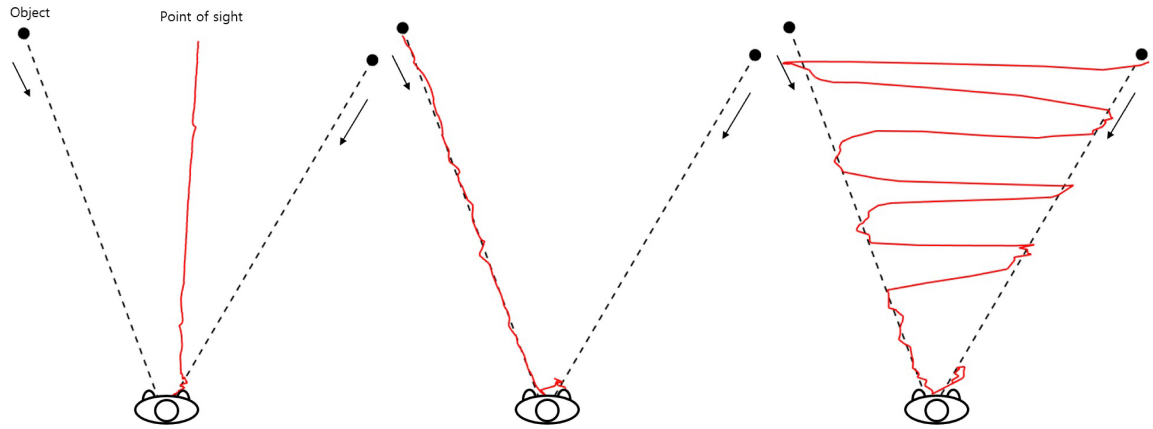

(a) Minimizing the distance sum (b) Yeo et al. [2012] (c) Our method

Fig. 4. Comparison of trajectories of the point of sight produced with different approaches. Dotted and red lines represent the trajectories of black balls and the character's point of sight, respectively. Use of the distance sum instead of the uncertainty sum makes the point of sight to be placed in the middle of the two balls at each time step (a). The method by Yeo et al. [2012] makes the character track the balls one after the other (b). Our method allows the character to automatically switch the point of sight between the two different balls (c).

the closest to the contact foot position extracted from the last optimal state trajectory at every frame of the time window. Task-dependent cost term c_{tsk} was formulated to make the contact foot as close as possible to the selected stone while preserving its global foot orientation to that of the reference motion as follows:

$$c_{tsk} = \|\mathbf{p}_{\text{foot}} - \mathbf{p}_{\text{cs}}\|^2 + \|\mathrm{W}_{\text{foot}} - \bar{\mathrm{W}}_{\text{foot}}\|^2,$$

where $\mathbf{p}_{\text{foot}}$ is the position of the contact foot, $\mathbf{p}_{\text{cs}}$ is the position of the closest stone to the foot, W_{foot} is the rotation matrix of the foot and $\bar{\mathrm{W}}_{\text{foot}}$ is the corresponding reference matrix. Note that c_{tsk} is applied only at contact frames. We set the value of w_{tsk} to 10.

In this experiment, the character first looked at the selected stone before moving the stepping foot to the stone and then looked at the next target stone once the foot began to move, as shown in Figure 5. These results are in line with the behavioral studies for humans in biomechanics [Matthis et al. 2015].

Balancing after a push

This experiment validates looking behaviors of a character standing near a wall modeled with multiple cube blocks when pushed with an external force. With the external force applied to the character for pushing, our system detects the moment, if there is any, at which the trajectory of the center of pressure (COP) crosses over the boundary of the supporting polygon for the feet of the character. At this moment, the character automatically moves a wall supporting hand toward the nearest estimated cube block of the wall while keeping the visual attention at the center of the block on which it places the hand. The character is supposed to make a transition from a standing pose to the hand support leaning motion that was captured without gaze behaviors and add the secondary behaviors to the motion on the fly. Finally, the character removes the hand from the wall once the character recovers its balance. Accordingly, it turns the head away from the wall. Indeed, our system performed this scenario very well as demonstrated in the supplementary video. Note that we did not use task-dependent cost term c_{tsk} for balancing.

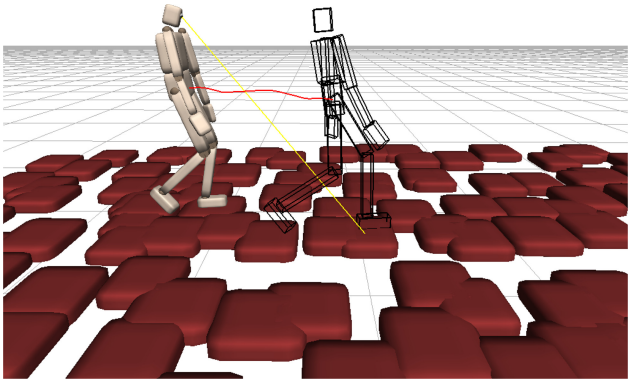

Fig. 5. Walking on stepping stones. The wireframe mesh represents the character pose at the last frame of the window, the red line is the root trajectory, and the yellow line is the line of sight. Once the foot begins to move, the character looks at the next target stone.

Moving obstacle avoidance

This experiment was performed on running motion to verify looking behaviors of a character while avoiding moving obstacles, of which the initial position and speed were chosen randomly. Among the perceived obstacles, the four closest ones to the character were selected for trajectory optimization at every time step to generate an optimal control policy based on their future belief states. Similar to Treuille et al. [2007], we formulated a task-dependent cost term c_{tsk}, which rapidly increased when the character moved close to any of the four obstacles as follows:

$$c_{tsk} = \sum_{j=0}^{N_{\text{obj}}-1} exp\left(\frac{-\|\mathbf{p}_{\text{root}} - \hat{\mathbf{p}}_j\|}{s_j^2}\right),$$

where $\mathbf{p}_{\text{root}}$ is the global root position, $\hat{\mathbf{p}}_j$ is the estimated position of the jth object, which is the first three components of $\hat{\mu}_j$ of each object's belief state, and s_j is the radius of jth object.

In general, the time window should be large enough to avoid moving obstacles while running (for example, more than 1s). In this experiment, we set the window size to 1.6s, which incurs heavy computational load. To achieve interactive performance, we divide the window into two sub-windows of size 1.0s and 0.6s, respectively. For the former sub-window, the system evaluates c_{tsk} following the root trajectory. For the latter sub-window, the system linearly extrapolates the root trajectory to evaluate c_{tsk} with the root velocity fixed at the value from the last frame of the former sub-window. We set the value of w_{tsk} to 1.0.

8 CONCLUSIONS

In this article, we presented an MPC framework with a visuomotor system that generates realistic gaze behaviors coupled with physics-based full-body motions. Based on human visual perception and full-body contact dynamics, we formulated an optimal motion control problem with a POMDP, which is known to be extremely difficult to solve. To address this issue, we approximated the POMDP as a deterministic belief MDP and solved it with DDP over a finite-horizon window while shifting the window along the time-axis. Guided by a reference skeletal motion without any *a priori* gaze information, our system produced realistic eye and head movements on top of full-body motions for various tasks, such as catching a thrown ball, walking on stepping stones, balancing after a push, and avoiding moving obstacles.

For interactive performance, we adopted deterministic DDP by taking the expectation of mean update μ_j' based on recent results [Erez and Smart 2012; Van Den Berg et al. 2012, 2017]. However, employing stochastic DDP [Theodorou et al. 2010] can be a viable alternative to maintain stochastic characteristics of POMDP. This choice would lead to a much complex and heavy system, since full-body dynamics is highly non-linear [Theodorou et al. 2010]. It can be a challenging and interesting future work to faithfully reproduce stochastic nature of POMDP in the trajectory optimization while maintaining its interactive performance and robustness.

Here are some limitations of our current visuomotor system: Due to the use of smoothed contact dynamics, an unexpected strong impact may cause slight penetration of a foot into the ground. Although being robust against external forces or environmental changes, our system inherits the limitation of an MPC framework in terms of generality. Based on the example-based approach, our system shows an unnatural result when the current state of a character is too different from that of the reference motion. Objects that move very quickly may induce instantaneous and quick head movements, which may also result in unnatural full-body motions. The complex non-linear equations for updating the covariance of a belief state incur a computational load that is too heavy to achieve truly real-time physics-based motion control.

Our work deals with simple external objects such as spheres and boxes. It would be an interesting research direction to consider complex objects such as living creatures. It would also be an interesting future direction to improve MPC for the complex visuomotor dynamics using deep reinforcement learning, which attracts ever-growing attention from the research community. Peripheral vision and eye blinking are the essential features of a human vision system, and their incorporation into the synthetic vision system would improve the naturalness of resulting gaze behaviors. We have modeled our vision system based on 3D object states composed of positions and velocities. To better mimic the human vision system, it would be a good future research direction to exploit visual observations such as binocular images for tracking moving objects near the character in the synthetic environment with a deeper understanding of the visual processing mechanism of human brains. We also note that the current version of our system does not take into account the uncertainty of the character's own uncertainty.

APPENDIX A

Based on Equations (1) and (2), our system performs the belief update through the prediction and update steps of a Kalman filter. Given current belief state $b(\mathbf{x}_j) = \mathcal{N}(\mu_j, \Sigma_j)$ for object j, the prediction step produces *a priori* belief estimate $\bar{b}(\mathbf{x}_j) = \mathcal{N}(\bar{\mu}_j, \bar{\Sigma}_j)$, where

$$\begin{aligned}
\bar{\mu}_j &= \mathrm{F}\mu_j + \mathbf{y} \text{ and} \\
\bar{\Sigma}_j &= \mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q} \\
&= \begin{bmatrix} (\sigma_j^{\mathrm{p}} + h\sigma_j^{\mathrm{pv}} + h(\sigma_j^{\mathrm{pv}} + h\sigma_j^{\mathrm{v}}) + \alpha)\mathrm{I} & (\sigma_j^{\mathrm{pv}} + h\sigma_j^{\mathrm{v}})\mathrm{I} \\ (\sigma_j^{\mathrm{pv}} + h\sigma_j^{\mathrm{v}})\mathrm{I} & (\sigma_j^{\mathrm{v}} + \alpha)\mathrm{I} \end{bmatrix} \\
&= \begin{bmatrix} \bar{\sigma}_j^{\mathrm{p}}\mathrm{I} & \bar{\sigma}_j^{\mathrm{pv}}\mathrm{I} \\ \bar{\sigma}_j^{\mathrm{pv}}\mathrm{I} & \bar{\sigma}_j^{\mathrm{v}}\mathrm{I} \end{bmatrix}
\end{aligned}$$

Combining observation $\mathbf{z}_j$ at the current time step with *a priori* belief estimate predicted above, the update step leads to *a posteriori* belief estimate, that is, belief update $b'(\mathbf{x}_j) = \mathcal{N}(\mu_j', \Sigma_j')$, where

$$\mu_j' = \bar{\mu}_j + \mathrm{K}_j\mathbf{m}_j \text{ and } \Sigma_j' = (\mathrm{I} - \mathrm{K}_j\mathrm{H})\bar{\Sigma}_j.$$

In the above equations, $\mathbf{m}_j$ and K_j denote the observation residual and the Kalman gain for each object, respectively, which are derived as follows:

$$\begin{aligned}
\mathbf{m}_j &= \mathbf{z}_j - \mathrm{H}\bar{\mu}_j = \mathbf{z}_j - \mathrm{H}(\mathrm{F}\mu_j + \mathbf{y}), \\
\mathrm{K}_j &= \bar{\Sigma}_j\mathrm{H}^{\mathrm{T}}\mathrm{S}_j^{-1},
\end{aligned}$$

where S_j is the residual covariance:

$$\mathrm{S}_j = \mathrm{H}\bar{\Sigma}_j\mathrm{H}^{\mathrm{T}} + \mathrm{R}_j = \mathrm{H}(\mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q})\mathrm{H}^{\mathrm{T}} + \mathrm{R}_j.$$

Therefore,

$$\begin{aligned}
\mu_j' &= \bar{\mu}_j + \mathrm{K}_j\mathbf{m}_j = \mathrm{F}\mu_j + \mathbf{y} + \mathrm{K}_j(\mathbf{z}_j - \mathrm{H}(\mathrm{F}\mu_j + \mathbf{y})), \\
\Sigma_j' &= (\mathrm{I} - \mathrm{K}_j\mathrm{H})\bar{\Sigma}_j = (\mathrm{I} - \mathrm{K}_j\mathrm{H})(\mathrm{F}\Sigma_j\mathrm{F}^{\mathrm{T}} + \mathrm{Q}).
\end{aligned}$$

APPENDIX B

Suppose that there is a single object in the environment. Then, the character would keep looking at the object. According to Equation (1), as the point of sight approaches to object j, R_j becomes a zero matrix, since r_j tends to a zero. Thus, we can derive the limit

of Kalman gain K_j as shown below:

$$\begin{aligned}\lim_{r_j\to 0}\mathrm{K}_j &= \lim_{r_j\to 0}\bar{\Sigma}_j\mathrm{H}^{\mathrm{T}}(\mathrm{H}\bar{\Sigma}_j\mathrm{H}^{\mathrm{T}}+\mathrm{R}_j)^{-1}\\ &= \lim_{r_j\to 0}\left(\frac{1}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\right)\begin{bmatrix}\bar{\sigma}_j^{\mathrm{p}}\mathrm{I}\\ \bar{\sigma}_j^{\mathrm{pv}}\mathrm{I}\end{bmatrix}\\ &= \begin{bmatrix}\mathrm{I}\\ \frac{\bar{\sigma}_j^{\mathrm{pv}}}{\bar{\sigma}_j^{\mathrm{p}}}\mathrm{I}\end{bmatrix}.\end{aligned}$$

Using the result above, we can derive the limit of μ_j' in Equation (3) as follows:

$$\begin{aligned}\lim_{r_j\to 0}\mu_j' &= \lim_{r_j\to 0}(\bar{\mu}_j+\mathrm{K}_j(\mathrm{z}_j-\mathrm{H}\bar{\mu}_j)\\ &= \bar{\mu}_j+\begin{bmatrix}\mathrm{I}\\ \frac{\bar{\sigma}_j^{\mathrm{pv}}}{\bar{\sigma}_j^{\mathrm{p}}}\mathrm{I}\end{bmatrix}(\mathrm{z}_j-\mathrm{H}\bar{\mu}_j)\\ &= \begin{bmatrix}\mathrm{p}_j\\ \bar{\mu}_j^{\mathrm{v}}+\frac{\bar{\sigma}_j^{\mathrm{pv}}}{\bar{\sigma}_j^{\mathrm{p}}}(\mathrm{p}_j-\bar{\mu}_j^{\mathrm{p}})\end{bmatrix}.\end{aligned}$$

As $r_j\to 0$, the estimated state approaches the true state in terms of the position but not in terms of the velocity. This means that we only have a perfect observation on the (true) object position as $r_j\to 0$. This is in line with our assumption made in Section 4.3 that humans perceive the positional information of a moving object to estimate its instantaneous velocity.

Σ_j' in Equation (4) is also reduced to the following equation,

$$\begin{aligned}\lim_{r_j\to 0}\Sigma_j' &= \lim_{r_j\to 0}(\mathrm{I}-\mathrm{K}_j\mathrm{H})\bar{\Sigma}_j\\ &= \lim_{r_j\to 0}\begin{bmatrix}\left(\frac{\bar{\sigma}_j^{\mathrm{p}}r_j}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\right)\mathrm{I} & \left(\frac{\bar{\sigma}_j^{\mathrm{pv}}r_j}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\right)\mathrm{I}\\ \left(\frac{\bar{\sigma}_j^{\mathrm{pv}}r_j}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\right)\mathrm{I} & \left(\bar{\sigma}_j^{\mathrm{v}}-\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\right)\mathrm{I}\end{bmatrix}\\ &= \begin{bmatrix}0 & 0\\ 0 & \left(\bar{\sigma}_j^{\mathrm{v}}-\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}}\right)\mathrm{I}\end{bmatrix}.\end{aligned}$$

Note that every submatrix tends to a zero as $r_j\to 0$ except for the velocity covariance matrix (the lower right submatrix of the last matrix). The norm of the velocity covariance matrix also takes on a minimum value as $r_j\to 0$, since $\bar{\sigma}_j^{\mathrm{v}}>\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}+r_j}$ for a positive constant α as shown below:

$$\begin{aligned}&\bar{\sigma}_j^{\mathrm{v}}-\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}+r_j}\\ &= (\sigma_j^{\mathrm{v}}+\alpha)-\frac{(\sigma_j^{\mathrm{pv}}+h\sigma_j^{\mathrm{v}})^2}{(\sigma_j^{\mathrm{p}}+h\sigma_j^{\mathrm{pv}}+h(\sigma_j^{\mathrm{pv}}+h\sigma_j^{\mathrm{v}})+\alpha)+r_j}\\ &= \alpha+\frac{\sigma_j^{\mathrm{v}}(\alpha+r_j)}{(\sigma_j^{\mathrm{p}}+h\sigma_j^{\mathrm{pv}}+h(\sigma_j^{\mathrm{pv}}+h\sigma_j^{\mathrm{v}})+\alpha)+r_j}>0\end{aligned}$$

Therefore, $\bar{\sigma}_j^{\mathrm{v}}>\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}+r_j}$.

In the third step of the above derivation, we make use of the fact that $\sigma_j^{\mathrm{p}}\sigma_j^{\mathrm{v}}=\sigma_j^{\mathrm{pv}}$ from the Cauchy-Schwarz inequality [Shi et al. 2007], because the object velocity has a linear relationship with the object position. Note that $\bar{\sigma}_j^{\mathrm{v}}-\frac{\bar{\sigma}_j^{\mathrm{pv}^2}}{\bar{\sigma}_j^{\mathrm{p}}}$ goes to a zero for a perfect process model, in which α has a zero value. Hence, when the character keeps looking at an object, our vision system tends to make the observed object state coincident with its true position with a minimal uncertainty.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their invaluable comments, Bumki Kim and Kyungmin Cho for modeling the character and rendering the scenes, and Seokpyo Hong for the helpful discussion about the experiments.

REFERENCES

Mazen Al Borno, Martin De Lasa, and Aaron Hertzmann. 2013. Trajectory optimization for full-body movements with complex contacts. *IEEE Trans. Vis. Comput. Graph.* 19, 8 (2013), 1405–1414.

George A. Alvarez and Steven L. Franconeri. 2007. How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism. *J. Vis.* 7, 13 (2007), 14–14.

Chris L. Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B. Tenenbaum. 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* 1, 4 (2017), 0064.

Boris Belousov, Gerhard Neumann, Constantin A. Rothkopf, and Jan R. Peters. 2016. Catching heuristics are optimal control policies. In *Advances in Neural Information Processing Systems*. 1426–1434.

Zhigang Deng, John P. Lewis, and Ulrich Neumann. 2005. Automated eye motion using texture synthesis. *IEEE Comput. Graph. Appl.* 25, 2 (2005), 24–30.

Noel E. Du Toit and Joel W. Burdick. 2012. Robot motion planning in dynamic, uncertain environments. *IEEE Trans. Robot.* 28, 1 (2012), 101–115.

Tom Erez and William D. Smart. 2009. Coupling perception and action using minimax optimal control. In *Proceedings of the 2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*. IEEE, 58–65.

Tom Erez and William D. Smart. 2012. A scalable method for solving high-dimensional continuous POMDPs using local approximation. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*. AUAI Press.

Tom Erez, Julian J. Tramper, William D. Smart, and Stan C. A. M. Gielen. 2011. A POMDP model of eye-hand coordination. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*.

Anthony C. Fang and Nancy S. Pollard. 2003. Efficient synthesis of physically valid human motion. *ACM Trans. Graph.* 22, 3 (2003), 417–426.

Roy Featherstone. 2008. *Rigid Body Dynamics Algorithms*. Vol. 49. Springer, Berlin.

Brian Guenter. 2007. Efficient symbolic differentiation for graphics applications. *ACM Trans. Graph.* 26, 3 (2007), 108.

Perttu Hämäläinen, Sebastian Eriksson, Esa Tanskanen, Ville Kyrki, and Jaakko Lehtinen. 2014. Online motion synthesis using sequential Monte Carlo. *ACM Trans. Graph.* 33, 4, Article 51 (July 2014), 12 pages. DOI: https://doi.org/10.1145/2601097.2601218

Perttu Hämäläinen, Joose Rajamäki, and C. Karen Liu. 2015. Online control of simulated humanoids using particle belief propagation. *ACM Trans. Graph.* 34, 4 (2015), 81.

Daseong Han, Haegwang Eom, Junyong Noh, and Joseph Shin. 2016. Data-guided model predictive control based on smoothed contact dynamics. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 533–543.

Daseong Han, Junyong Noh, Xiaogang Jin, and Joseph Shin. 2014. On-line real-time physics-based predictive motion control with balance recovery. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 245–254.

David B. Henson et al. 2000. *Visual Fields*. Vol. 457. Butterworth-Heinemann.

Laurent Itti. 2006. Quantitative modelling of perceptual salience at human eye position. *Vis. Cogn.* 14, 4–8 (2006), 959–984.

David Jacobson and David Mayne. 1970. *Differential Dynamic Programming*. American Elsevier, New York.

Sumit Jain, Yuting Ye, and C. Karen Liu. 2009. Optimization-based interactive motion synthesis. *ACM Trans. Graph.* 28, 1 (2009), 10.

Vikash Kumar, Yuval Tassa, Tom Erez, and Emanuel Todorov. 2014. Real-time behaviour synthesis for dynamic hand-manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'14)*. IEEE, 6808–6815.

Brent J. Lance and Stacy C. Marsella. 2010. *The Expressive Gaze Model: Using Gaze to Express Emotion.* Technical Report. DTIC Document.
Sooha Park Lee, Jeremy B. Badler, and Norman I. Badler. 2002. Eyes alive. *ACM Trans. Graph.* 21, 3 (July 2002), 637–644. DOI: https://doi.org/10.1145/566654.566629
C. Karen Liu, Aaron Hertzmann, and Zoran Popović. 2005. Learning physics-based motion style with nonlinear inverse optimization. *ACM Trans. Graph* 24, 3 (July 2005), 1071–1081. DOI: https://doi.org/10.1145/1073204.1073314
Libin Liu, KangKang Yin, and Baining Guo. 2015. Improving sampling-based motion control. *Comput. Graph. Forum* 34, 2 (2015).
Libin Liu, KangKang Yin, Michiel van de Panne, Tianjia Shao, and Weiwei Xu. 2010. Sampling-based contact-rich motion control. *ACM Trans. Graph.* 29, 4 (2010), 128.
Omid Madani, Steve Hanks, and Anne Condon. 2003. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artif. Intell.* 147, 1 (2003), 5–34.
Jonathan Samir Matthis, Sean L. Barton, and Brett R. Fajen. 2015. The biomechanics of walking shape the use of visual information during locomotion over complex terrain. *J. Vis.* 15, 3 (2015), 10–10.
Joseph McIntyre, M. Zago, A. Berthoz, and F. Lacquaniti. 2001. Does the brain model Newton's laws? *Nat. Neurosci.* 4, 7 (2001), 693–694.
Craig H. Meyer, Adrian G. Lasker, and David A. Robinson. 1985. The upper limit of human smooth pursuit velocity. *Vis. Res.* 25, 4 (1985), 561–563.
Igor Mordatch, Emanuel Todorov, and Zoran Popović. 2012. Discovery of complex behaviors through contact-invariant optimization. *ACM Trans. Graph.* 31, 4 (2012), 43.
Masaki Nakada, Tao Zhou, Honglin Chen, Tomer Weiss, and Demetri Terzopoulos. 2018. Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Trans. Graph.* 37, 4 (2018), 56.
Debanga R. Neog, João L. Cardoso, Anurag Ranjan, and Dinesh K. Pai. 2016. Interactive gaze driven animation of the eye region. In *Proceedings of the 21st International Conference on Web3D Technology (Web3D'16)*. ACM, New York, NY, 51–59. DOI: https://doi.org/10.1145/2945292.2945298
Christos H. Papadimitriou and John N. Tsitsiklis. 1987. The complexity of Markov decision processes. *Math. Operat. Res.* 12, 3 (1987), 441–450.
Tomislav Pejsa, Sean Andrist, Michael Gleicher, and Bilge Mutlu. 2015. Gaze and attention management for embodied conversational agents. *ACM Trans. Interact. Intell. Syst.* 5, 1 (2015), 3.
Tomislav Pejsa, Daniel Rakita, Bilge Mutlu, and Michael Gleicher. 2016. Authoring directed gaze for full-body motion capture. *ACM Trans. Graph.* 35, 6 (2016), 161.
Christopher Peters and Adam Qureshi. 2010. A head movement propensity model for animating gaze shifts and blinks of virtual characters. *Comput. Graph.* 34, 6 (2010), 677–687.
Joelle Pineau and Geoffrey Gordon. 2007. POMDP planning for robust robot control. In *Robotics Research*, S. Thrun, R. Brooks, and H. Durrant-Whyte (Eds.). Springer Tracts in Advanced Robotics, Vol. 28. Springer, Berlin, Heidelberg, 69–82.
Robert Platt, Jr, Russ Tedrake, Leslie Kaelbling, and Tomas Lozano-Perez. 2010. Belief space planning assuming maximum likelihood observations. In *Robotics Science and Systems Conference (RSS'10)*.
Zoran Popović and Andrew Witkin. 1999. Physically based motion transformation. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., 11–20.
Rajesh P. N. Rao and Dana H. Ballard. 1997. Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neur. Comput.* 9, 4 (1997), 721–763.
D. A. Robinson. 1965. The mechanics of human smooth pursuit eye movement. *J. Physiol.* 180, 3 (1965), 569.
Kerstin Ruhland, Christopher E. Peters, Sean Andrist, Jeremy B. Badler, Norman I. Badler, Michael Gleicher, Bilge Mutlu, and Rachel McDonnell. 2015. A review of eye gaze in virtual agents, social robotics and HCI: Behaviour generation, user interaction and perception. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 299–326.
Alp Sardag and H. Levent Akin. 2006. Kalman based finite state controller for partially observable domains. *Int. J. Adv. Robot. Syst.* 3, 4 (2006), 45.
Brian J. Scholl and Yaoda Xu. 2001. The magical number 4 in vision. *Behav. Brain Sci.* 24, 01 (2001), 145–146.
Raymond A. Serway, Robert J. Beichner, and John W. Jewett. 2018. Physics for scientists and engineers with modern physics. Cengage Learning.
Ling Shi, Karl Henrik Johansson, and Richard M. Murray. 2007. Kalman filtering with uncertain process and measurement noise covariances with application to state estimation in sensor networks. In *Proceedings of the 2007 IEEE International Conference on Control Applications*. IEEE, 1031–1036.
Y. Shin, H. W. Lim, M. H. Kang, M. Seong, H. Cho, and J. H. Kim. 2016. Normal range of eye movement and its relationship to age. *Acta Ophthalmol.* 94, S256 (2016).
Edward Jay Sondik. 1971. *The Optimal Control of Partially Observable Markov Processes.* Technical Report. DTIC Document.
Matthijs T. J. Spaan. 2012. Partially observable Markov decision processes. In *Reinforcement Learning*. Springer, 387–414.
Robert F. Stengel. 1994. *Optimal Control and Estimation.* Courier Corporation.
Yuval Tassa, Tom Erez, and Emanuel Todorov. 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4906–4913.
Evangelos Theodorou, Yuval Tassa, and Emo Todorov. 2010. Stochastic differential dynamic programming. In *Proceedings of the American Control Conference (ACC'10)*. IEEE, 1125–1132.
Sebastian Thrun. 2000. Monte Carlo pomdps. In *Advances in Neural Information Processing Systems*. 1064–1070.
Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics.* MIT Press.
E. Todorov. 2014. Convex and analytically invertible dynamics with contacts and constraints: Theory and implementation in MuJoCo. In *Proceedings of the IEEE Conference on Robotics and Automation (ICRA'14)*.
Adrien Treuille, Yongjoon Lee, and Zoran Popović. 2007. Near-optimal character animation with continuous control. *ACM Trans. Graph.* 26, 3, Article 7 (July 2007). DOI: https://doi.org/10.1145/1276377.1276386
Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. 2012. Motion planning under uncertainty using iterative local optimization in belief space. *Int. J. Robot. Res.* 31, 11 (2012), 1263–1278.
Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. 2017. Motion planning under uncertainty using differential dynamic programming in belief space. In *Robotics Research*. Springer, 473–490.
Greg Welch and Gary Bishop. 2006. An introduction to the kalman filter. Department of Computer Science, University of North Carolina. (unpublished).
Andrew Witkin and Michael Kass. 1988. Spacetime constraints. *ACM Siggraph Comput. Graph.* 22, 4 (1988), 159–168.
Yuting Ye and C. Karen Liu. 2010. Optimal feedback control for character animation using an abstract model. *ACM Trans. Graph.* 29, 4, Article 74 (July 2010), 9 pages. DOI: https://doi.org/10.1145/1778765.1778811
Sang Hoon Yeo, Martin Lesmana, Debanga R. Neog, and Dinesh K. Pai. 2012. Eyecatch: Simulating visuomotor coordination for object interception. *ACM Trans. Graph.* 31, 4 (2012), 42.
R. J. Leigh and D. S. Zee. 2015. The neurology of eye movements. OUP USA.

Received November 2018; revised August 2019; accepted August 2019