Natural Sciences and Engineering Research Council of Canada
Conseil de recherches en sciences naturelles et en génie du Canada

1507

Institutional Identifier

System-ID (for NSERC use only)
125523743

FORM 101
Application for a Grant
PART I

Date
2009/10/26

Family name of applicant van de Panne	Given name Michiel	Initial(s) of all given names M	Personal identification no. (PIN) **Valid** 103212

Institution that will administer the grant
British Columbia

Language of application
☒ English ☐ French

Time (in hours per month) to be devoted to the proposed research / activity
50

Type of grant applied for
Discovery Grants - Individual

For Strategic Projects, indicate the Target Area and the Research Topic; for Strategic Networks and Strategic Workshops indicate the Target Area.

Title of proposal
Scalable Physics-based Modeling of Skilled Movement

Provide a maximum of 10 key words that describe this proposal. Use commas to separate them.
character animation, computer graphics, physics-based simulation, motor control, humanoid robotics, motion planning

Research subject code(s)		Area of application code(s)	
Primary	Secondary	Primary	Secondary
2707	2600	801	1200

CERTIFICATION/REQUIREMENTS

If this proposal involves any of the following, check the box(es) and submit the protocol to the university or college's certification committee.

Research involving : Humans ☐ Human pluripotent stem cells ☐ Animals ☐ Biohazards ☐

Does any phase of the research described in this proposal a) take place outside an office or laboratory, or b) involve an undertaking as described in Part 1 of Appendix B?

☒ NO ☐ If YES to either question a) or b) – Appendices A and B must be completed

TOTAL AMOUNT REQUESTED FROM NSERC

Year 1	Year 2	Year 3	Year 4	Year 5
105,600	105,600	105,600	105,600	105,600

SIGNATURES (Refer to instructions "What do signatures mean?")

It is agreed that the general conditions governing grants as outlined in the NSERC *Program Guide for Professors* apply to any grant made pursuant to this application and are hereby accepted by the applicant and the applicant's employing institution.

Applicant

Head of department

Applicant's department, institution, tel. and fax nos., and e-mail
Computer Science
British Columbia
Tel.: (604) 822-8737
FAX: (604) 822-4231
van@cs.ubc.ca

Dean of faculty

President of institution
(or representative)

The information collected on this form and appendices will be stored in the Personal Information Bank for the appropriate program.
PROTECTED WHEN COMPLETED

Version française disponible

Canada

Personal identification no. (PIN) **Valid** 103212	Family name of applicant van de Panne

SUMMARY OF PROPOSAL FOR PUBLIC RELEASE (Use plain language.)

This plain language summary will be available to the public if your proposal is funded. Although it is not mandatory, you may choose to include your business telephone number and/or your e-mail address to facilitate contact with the public and the media about your research.

Business telephone no. (optional):

E-mail address (optional): van@cs.ubc.ca

Humans and animals move through their world and interact with their environment in robust, graceful, and highly skilled ways. We do not yet know how to design realistic animated characters, biomechanical human simulations, and robots that can mimic this degree of skill. However, recent developments suggest that significant progess is being made towards mimicing human-like dexterity in simulations, or, in some cases, with robots. In this proposal, we use robust simulations of bipedal walking as a point of departure for developing richer movement skill sets, including highly agile walking and running, natural movement through urban environments, and dexterous hand manipulation. We target next generation of computer animation software as our primary application, but we expect that the same models will also be useful for understanding and correcting pathological gaits, the design of walking prostheses, obtaining a deeper understanding of human motor control, and creating more skillful robots.

Other Language Version of Summary (optional).

Personal identification no. (PIN)	Family name of applicant
Valid 103212	van de Panne

Before completing this section, **read the instructions** and consult the *Use of Grant Funds* section of the NSERC Program Guide for Professors concerning the eligibility of expenditures for the direct costs of research and the regulations governing the use of grant funds.

TOTAL PROPOSED EXPENDITURES (Include cash expenditures only)

	Year 1	**Year 2**	**Year 3**	**Year 4**	**Year 5**
1) Salaries and benefits					
a) Students	64,500	64,500	64,500	64,500	64,500
b) Postdoctoral fellows	9,000	9,000	9,000	9,000	9,000
c) Technical/professional assistants	7,100	7,100	7,100	7,100	7,100
d)	0	0	0	0	0
2) Equipment or facility					
a) Purchase or rental	8,000	8,000	8,000	8,000	8,000
b) Operation and maintenance costs	0	0	0	0	0
c) User fees	0	0	0	0	0
3) Materials and supplies	3,000	3,000	3,000	3,000	3,000
4) Travel					
a) Conferences	12,000	12,000	12,000	12,000	12,000
b) Field work	0	0	0	0	0
c) Collaboration/consultation	2,000	2,000	2,000	2,000	2,000
5) Dissemination costs					
a) Publication costs	0	0	0	0	0
b)	0	0	0	0	0
6) Other (specify)					
a)	0	0	0	0	0
b)	0	0	0	0	0
TOTAL PROPOSED EXPENDITURES	105,600	105,600	105,600	105,600	105,600
Total cash contribution from industry (if applicable)					
Total cash contribution from university (if applicable)					
Total cash contribution from other sources (if applicable)	0	0	0	0	0
TOTAL AMOUNT REQUESTED FROM NSERC (transfer to page 1)	105,600	105,600	105,600	105,600	105,600

Budget Justification

1a Students

I am budgeting for $16,500 per year for M.Sc. students and $19,000 per year for PhD students, as per NSERC rules for maximum salaries from grant funds. I expect to fund undergraduate internships at an average of $5,000 per year. Postdoctoral fellows cost $45,000 per year (including benefits). I am requesting $64,500 in order to support 2 PhD, 1 MSc, and 2 undergraduate students per year. I expect to host a postdoctoral fellow for at least one of the five years of the grant. The exact year in which I expect to hire the postdoctoral fellow is not known as I expect to do so in an opportunistic fashion in order to obtain the best possible candidate. I have therefore distributed the cost equally over the five years (5 × $9,000).

1c Technical/Professional Assistants

The UBC Department of Computer Science charges the direct costs of technical support to research grants. This includes installation and support of equipment; technical support for researchers; print, file and network servers; printing and copying; and similar direct costs. The current charge is 2%, and thus $2,100 is requested for these research costs.

$5,000 is requested as my contribution towards a 10% share of a future lab manager (or equivalent technical support arranged in another fashion) to support our group of 10 graphics and HCI faculty. Tasks to be managed include the maintenance of experimental apparatus (including motion capture system, 3D scanner, and other specialized hardware, software licensing, etc.) and the management of resource-related issues associated with running the graphics, visualization, and HCI lab.

2a Purchase or Rental

$8,000 per annum provides for the purchase and upgrades of PCs, printers, and networking equipment for the research needs of the graduate and undergraduate students working under my supervision. This includes two PCs (2 × $2,000), $2,000 for software and software licenses, and $2,000 for printers, networking, and other miscellaneous hardware.

4a Conference Travel

$12,000 will provide for 3 conference-related trips per year for myself (SIGGRAPH, Symposium on Computer Animation, Graphics Interface, ICRA, as well as program committee meetings) and an average of one conference trip per year for each of 5 students and postdoctoral fellows. Conference attendance costs are estimated to be $1,500.

4c Collaboration travel

$2,000 will provide for two trips per year, to be shared between myself and my students, in order to develop contacts and enable collaborations that lie outside of computer graphics and animation, i.e., robotics, biomechanics, and predictive human modeling for healthcare applications.

Relationship to Other Research Support

The NSERC Discovery Grant is my principal source of research funding for graduate student support, conference travel, and basic equipment purchases (non-specialized PCs).

MITACS (ongoing)
An ongoing MITACS grant provided for $4–7k per annum, is used in partial support of students and equipment, e.g., Tablet PCs, for pen-based interaction applications, which are an ongoing research interest. The renewal of this grant remains dependent on our ability to find matching funds from industry.

Canada Research Chair (ongoing)
My Tier 2 CRC research funding goes to UBC and does not provide me with any research support.

GRAND NCE application (submitted)
I am one of 50 network investigators for an NCE application, entitled “GRAND”. My share would be approximately 2%, or $93k per year. I expect to spend $33k on funding collaborating researchers, who are not listed as network investigators. The remaining $60k, if awarded, will go towards funding a post-doc ($45 k), and another M.Sc. student focussed on physics-based character for games, in keeping with the ”Games, Animation, and New Media” topic of the NCE. My NSERC Discovery funds target the development of more abstract principles for motion control, including open theoretical issues (learning for motor control, more realistic sensori-motor control models) and applications to humanoid robotics, as well as other topics of interest in computer graphics that will be investigated as time and opportunity allows, e.g., geometric modeling, rendering.

1 Introduction

The ability of humans and animals to move with intelligence and grace through their environment remains unsurpassed. It is clear that the state of the art in computer graphics and robotics still falls well short of what we can observe in nature.

Computer animation offers a fertile playground for simulation-based modeling of skilled motion and sensori-motor control. It allows for the freedom to experiment with a wide variety of representations without being unduly concerned with the limitations of robotic hardware or precise measures of the biomechanical fidelity of the final motions, while still demanding results that are at least visually plausible. Also, computer animation respects and encourages the development of solutions that may lie anywhere along the automation spectrum, ranging from solutions requiring heavy manual (or data-driven) specification at one end, to solutions that are highly automated by deriving motions from a compact set of governing principles. This allows for the development of partial solutions at many points on this spectrum that can be used as stepping stones to more complete solutions. While not without its own limitations and tool-induced biases, a simulation-based methodology allows for convenient and flexible exploration of a large variety of architectures related to the control of motion. Physics-based models of locomotion are of particular interest to me, as locomotion plays a role in so many of the motions that we wish to be able to model and understand.

2 Progress

The lack of robust and flexible control of locomotion has long been a sticking point for physics-based character animation. However, over the past four years, work by a number of research groups [16, 9, 8] and my own group [4, 5, 12, 13] has resulted in exciting advances.

Our group has made particularly exciting progress on a number of areas in the past three years, including: simple and robust balance strategies [4]; novel optimization techniques for creating new skills [5, 12]; and introducing methods for environment-aware and task-aware planning of real-time physics-based locomotion [12, 13]. Our most recent work [13] contains two prototype games that have simple game play but that demonstrate the great potential of characters with physics-based skills.

We have distilled a number of key lessons from our recent work, which we intend to build on: **Abstraction:** Good abstractions of tasks and actions matter. For example, the actions in [RL] represent complete steps and consist of making a discrete choice among a small set of controllers, e.g., [4, 13]. **Focussed Policies:** Resources for representing and computing control policies should be focussed on restricted regions of state space that actually matter, i.e., that are commonly encountered, e.g., [12, 9, 13]. **Incremental learning:** New motion skills can be efficiently learned by adapting existing skills, e.g., [5, 16]. **Tuple-based Dynamics:** Sets of (state,action,new state) tuples can provide a compact-yet-flexible non-parametric representation of state dynamics [12, 13]. **Simple Feedback Laws:** Motions can often either be described in terms of simple objectives [8] or be stabilised by simple feedback laws [4, 10].

3 Objectives and Methodology

The over-arching objectives of this proposal are to develop more *scalable* solutions for developing simulated characters that can move with skill and grace. Scalable refers to stretching the existing methods in any number of dimensions, such as: controllers that can make *any* body shape walk; extending control methods to less realistic (supernatural) and more realistic (biomechanical) motions; allowing large online communities to create and share animation controller artifacts; simulating crowds of physics-based characters; and learning to work with more arbitrary sensori-motor soup instead of well-structured state and action variables.

We further break the proposed research into two groups. First, we focus on a specific problem, namely skilled biped locomotion. Second, we examine more abstract ideas related to learning and architectures for sensori-motor control. Because of space constraints, we focus on the ideas, methodology, significance, and evaluation of each project rather than the (often voluminous) previous work.

3.1 Skilled Biped Locomotion

Can we develop real-time physics-based simulations with biomechanically plausible actuation?

Current real-time physics-based locomotion simulations are robust and real time, but use highly abstracted torque actuation. More realistic actuation [3] requires introducing musculotendon geometry and muscle activation dynamics into the simulation, as well as computing muscle activations instead of joint torques. While the inversion of muscle dynamics is an ill-posed problem [14], the use of the musculoskeletal system also has the potential to simplify the control problem [2]. We are not aware of demonstrations of real-time biomechanically-plausible 3D walking simulations that have closed-loop balance feedback. We propose to use continuation methods [5] to progressively make the change from torque-based control to muscle activation control, while optimizing for robustness and energy usage. After doing initial experiments with a simple Hill-type muscles and simple line-of-force models, our next step will be to qualitatively evaluate the resulting behavior for differences in robustness and style. In the longer term, we foresee a collaboration with a biomechanics expert with a shared interest in real-time biomechanical models.

Can we create exaggerated, non-physical motions using physical simulations?

Going in the other direction, computer animation allows for the creation of imaginary creatures and imaginary worlds. This research direction proposes to allow for the addition of controlled external forces in order to achieve supernatural motions, e.g., giant jumps and leaps by a character. In this framework, physics and forces become a common currency that characters can use to interact with each other and the surrounding world. The idea is to reap the benefits of physical simulation e.g., collision response and momentum effects, while allowing for richer and more expressive classes of motion. We shall begin with manual experimentation in order to further develop our intuition about the problem. With the acquired knowledge about where and how to consider applying external forces and torques, we shall then move on to investigating the automated develoment of controllers that are capable of producing parameterized families of supernatural motions. Once we have developed a sufficient suite of supernatural controllers, the system will be iteratively evaluated and refined using feedback from animators.

Can physics-based walking skills be parameterized to support widely varying character scale, proportions, and mass distribution?

An answer to this problem would allow for any biped character to immediately come alive using physics-based simulation after being created. We will use dimensional scaling to first perform partial normalization and then use a regression-based estimator to predict the required controller parameters, in the face of remaining non-linearities. The samples required for the regression will first be computed on a sparse, regular grid as a sanity check. In the longer term we expect to apply an active learning approach for sample placement. If the variation of character parameters is kept small enough, we are guaranteed success. Thus for this project success will be defined by being able to model a parameterized controller across a large enough range of character parameters so as to be interesting and useful.

How can we develop task-specific models of natural footwork?

Many everyday tasks make use of the feet in ways that are neither standing nor walking, i.e., small reorientations or repositioning actions of the feet, small sideways steps, a left-and-right step to move to a new standing location, and so forth. Such motions are needed for natural and fluid-looking behaviors and cannot be produced by current techniques. We believe that an enriched inverted pendulum models may provide a compact and highly-predictive dynamics model to guide the planning and control of such motions for physics-based characters. At the same time, we intend to develop statistical models of context-dependent footwork that occurs in natural tasks. Success for this project is defined by the production of motions that are difficult or impossible to distinguish from motion capture data of observed unscripted behavior.

Can physics-based characters be used for realistic crowd simulation?

Although significant progress has been made, the characters in current crowd simulations often still behave in a wooden manner and this is particularly the case when modeling dense crowds. We expect that physics-based characters that are balance aware and equipped with a flexible array of walking and stepping skills can be used to make significant improvements to the state of the art. This requires the development of 'crowd skills' for characters as well as learning how to scale current simulation techniques to a large number of characters. Possible approaches include the development of simulations that are compatible with large fixed time steps and developing a suite of models of varying complexity (dynamic and kinematic) that can be invoked as allowed by the available computational budget. Of particular interest is that the approximation methods should also preserve the emergent properties of the crowd. The system will be evaluated by comparison to existing crowd animation systems.

3.2 Learning Motion Skills

We now propose a number of more abstract problems where we speculate that good progress can be made in the 5 years covered by this proposal. We take the stance that learning is core to building scalable sensori-motor control solutions, and that success lies with having the right architectures, representations, and training sequences to facilitate the learning. We speculate that learning-by-doing is a more powerful than learning-by-watching, and therefore we favor embodied learning [11, 15] over imitation-based methods.

How can skills be learned from scratch?

We propose thinking of a complex skill as being the end product of a set of relevant training exercises, during which motions are repeatedly optimized and integrated. In the case of locomotion, we propose to begin with a single in-place stepping gait as a point of departure. Using a suite of optimization metrics, this can be adapted to do a variety of things, such as walking forwards and backwards, walking and then leaping, and running. Given these gaits, a next task can be to learn the best ways to transition between them and therefore also how they can be sequenced. A last level of learning can then learn the right action to apply in a given task context. In the past two years we have demonstrated considerable success in gait optimization [5] and task-based control [13]. This *skills de novo* project is inspired by the integration of these ideas, as well as the idealistic goal of seeing structured, skilled motion emerge as a result of adaptive task-based interactions with the environment. In the first two years we will work with simplified two dimensional systems, after which we will consider agile 3D locomotion.

Can lessons learned from locomotion be applied to dexterous hand manipulation, and vice-versa?

Locomotion is an instance of dexterous manipulation, where the legs and the environment serve to manipulate the body. While there has been significant progress on grasping, e.g., [6, 7], there has been remarkably less progress on truly dexterous manipulation [1]. We intend to apply a number of the core ideas listed in the introduction towards developing simple-but-dexterous hand manipulation strategies [10], and then investigate scaling towards more complex strategies. A specific starting point here is to experiment with a large variety of hand manipulations that are modulated in simple ways by visual and tactile feedback [10]. This includes developing and testing models of finger gaiting. We expect that longer term work in this direction will involve collaborations with Prof. Dinesh Pai and his sensorimotor lab at UBC.

How can good control policies emerge in a sensori-motor soup?

Many approaches to control assume clean, minimal-length descriptions of the states and actions, whereas biologically-plausible encodings of sensori-motor problems would have a panoply of sensory inputs and actuation outputs. The sensory encodings are likely to use sparse, distributed representations and to include redundant and irrelevant information. Developing a compact policy description in this setting involves the learning of significant structure in the sensory and action encodings. However, there is a chicken-and-egg problem: (a) learning good control policies first requires having a manageable model of the state-and-action dynamics, and (b) learning a meaningful model of the state-and-action dynamics requires a good control policy in action. This bootstrapping problem is core to sensori-motor learning and has been examined in simpler contexts [15, 11]. Our goal with this project is to examine it in the context of the much higher dimensional settings offered by locomotion and manipulation problems. The bootstrapping problem will be addressed by developing a set of tasks that are learned in an easiest-to-hardest order.

What should be optimized to achieve human-like motion?

This can be posed as an inverse reinforcement learning or apprenticeship learning problem. Given human motion sequences as input, the goal is to determine how to weight a set of reward terms such that it would best explain the observed data. It is often easier to demonstrate natural behavior than to explain the reasons behind the behavior and therefore apprenticeship learning has much to offer to computer animation. While initial progress has been demonstrated for inverse learning for physics-based motion styles, we will focus on building models of human navigation and task-specific body motions. We will begin by gaining familiarity with current apprenticeship learning techniques on a well defined toy problem. We will then scale to working with footprint prediction problems, where the system is required to predict human-like footfall patterns based upon the environment context.

How can we learn the best feedback structures for stabilizing motions?

Both locomotion and object manipulation can be thought of as consisting of an open-loop execution plan that is then augmented by a set of fast, task-specific feedback loops. The feedback uses the sensory information to perform the many small adaptations that are needed for an otherwise-unstable motion to succeed. Early work on locomotion by Raibert and colleagues and our related work [4] has this general structure, where the inputs, outputs, and gains of the feedback loop are manually specified. In this project, the idea is to only provide limited knowledge regarding the structure of the feedback loop as a prior and to then learn the best sparse feedback structure and gains that can stabilise a given motion. We will apply a policy-search framework to this problem and use the first year to test it on locomotion problems where we know that such compact solutions exist. In subsequent years, we will apply it to a broader range of problems. Compact feedback structures are significant because they help achieve generalization of motor skills. As an example, constructing feedback strategies based on the motion of the center of mass helps greatly simplify the task of maintaining balance.

4 Training of Highly Qualified Personnel

My past research program has been highly successful at training HQP. Since 2004, I have supervised or cosupervised 17 MSc, PhD, and PDFs. Of these, 7 of them are now employed in the game industry in Vancouver, 4 are ongoing, 3 have found employment outside of Canada (associate researcher with Microsoft Research Asia, lecturer (faculty) at University College Dublin, and Microsoft Redmond), and 1 is completing a PhD program elsewhere after compleing his MSc with me.

Matching between HQP and the projects listed above will necessarily happen on an opportunistic basis, as I have done on past research projects. There needs to be an appropriate match between a students interests and skillset and those of the project. The current research proposal will train HQP in a variety of ways: reading groups for discussion and evaluation of the current state of the art; experience in collecting and processing motion data; launching a toolset for an online community to share animation artifacts; learning to establish a research program that aims high while taking achievable intermediate steps; writing and presenting papers; and networking with the relevant people in industry (games and film, where our group now has many connections) and academia. The proposed research also takes a broad view on the problem of motion control, and thus the students will develop skills in optimization, machine learning, and control, which are broadly used in many domains.

The HQP training from the proposed research will provide a skillset that is likely to be highly valued in the near future. Based on the pace of recent advances and on anecdotal comments I have observed from industry, it is quite likely that during the next five years we will see physics-based character animation enter the mainstream as a tool for animating motion in games, film, and scenario simulations. The same skill set is further equally valuable in other application areas, including robotics and biomechanics.

References

[1] A. Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on robotics and automation*, 16(6):652–662, 2000.

[2] R. Blickhan et al. Intelligence by mechanics. *Phil. Trans. of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365:199, 2007.

[3] F. Iida, J. Rummel, and A. Seyfarth. Bipedal walking and running with spring-like biarticular muscles. *Journal of Biomechanics*, 41(3):656–667, 2008.

[4] **K. Yin**, **K. Loken**, and M. van de Panne. SIMBICON: Simple biped locomotion control. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 26(3), 2007.

[5] **K. Yin**, **S. Coros**, **P. Beaudoin**, and Michiel van de Panne. Continuation methods for adapting simulated skills. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 27(3):81, 2008.

[6] Y. Li, J.L. Fu, and N.S. Pollard. Data-driven grasp synthesis using shape matching and task-based pruning. *IEEE Trans. on Visualization and Computer Graphics*, 13(4):732–747, 2007.

[7] C.K. Liu. Dextrous manipulation from a grasping pose. *ACM Trans. Graphics*, 28(3):59, 2009.

[8] A. Macchietto, V. Zordan, and C.R. Shelton. Momentum control for balance. *ACM Transactions on Graphics (Proc. SIGGRAPH*, 28(3):80, 2009.

[9] U. Muico, Y. Lee, J. Popović, and Z. Popović. Contact-aware nonlinear control of dynamic characters. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, page 81, 2009.

[10] A. Namiki and M. Ishikawa. Robotic catching using a direct mapping from visual information to motor command. In *IEEE Int’l Conf. on Robotics and Automation*, pages 2400–5, 2003.

[11] R. Pfeifer, M. Lungarella, and F. Iida. Self-organization, embodiment, and biologically inspired robotics. *Science*, 318:1088, 2007.

[12] **S. Coros**, **P. Beaudoin**, **K. Yin**, and M. van de Panne. Synthesis of constrained walking skills. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA)*, 27(5), 2008.

[13] **S. Coros**, **P. Beaudoin**, and M. van de Panne. Robust task-based control policies for physics-based characters. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA)*, 28(5), 2009. To appear.

[14] W. Tsang, K. Singh, and E. Fiume. Helping hand: an anatomically accurate inverse dynamics solution for unconstrained hand motion. In *Symp. on Computer animation*, page 328, 2005.

[15] P. Verschure, T. Voegtlin, and R. J. Douglas. Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425:620–624, 2003.

[16] J. M. Wang, D. J. Fleet, and A. Hertzmann. Optimizing walking controllers. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)*, 2009. to appear.