A Database Inference Controller for 3D Motion Capture Databases

Bhavani Thuraisingham, The University of Texas at Dallas, USA

Balakrishnan Prabhakaran, The University of Texas at Dallas, USA

Latifur Khan, The University of Texas at Dallas, USA

Kevin W. Hamlen, The University of Texas at Dallas, USA

ABSTRACT

This article presents a strategy for restructuring private human motion capture data to enforce access and inference controls within a relational database management system. Human 3D motion capture data is an important part of electronic health records for patients with motion-related diseases and symptoms. There are significant privacy concerns regarding the safe storage and dissemination of such data. Access controls traditionally applied to other forms of medical data (e.g., textual data) are not well suited to motion captures, which contain large quantities of data with complex interdependencies that divulge privacy-violating inferences. Encoding such motion data effectively within a relational database brings the large body of relational database security research to bear on this important problem.

Keywords: *access controls; computer science; healthcare; multimedia IS; relational database security*

1 INTRODUCTION

Human *motion analysis* is increasingly used by the medical community to diagnose, treat, and research a variety of conditions, including sports injuries (Fleisig et al., 1995), Parkinsonism (Morris et al., 2001), spinal cord injuries (Kaneoka et al., 1999), and conditions requiring prosthetics (Rietman et al., 2002). Data from the patient's motions are typically captured optically or from wearable sensors during a relevant exercise, such as treadmill walking or running (Lucareli et al., 2011), and stored in a 3D motion capture database for later analysis.

To address patients' privacy concerns regarding safe storage of electronic health records, access to each patient's motion data should ideally be controlled according to a *least privilege* principle (Saltzer, 1974)—that is, medical personnel should only have access to the portion of the data that is relevant to the condition they are treating. For example, a podiatrist treating a foot injury probably does not need access to fabric sensor data that records minute upper torso motions (Campbell et al., 2007). Withholding access privileges for unneeded portions of the data is important for mitigating the effects of attacks. For example, if an attacker who steals or forges the podiatrist's electronic credentials gains only limited (rather than full) access to the patient's medical records, the limitations serve to reduce the resulting privacy violation.

Unfortunately, traditional storage strategies for motion capture databases are not conducive to fine-grained access control because they typically encode motion data in monolithic files containing data from many or all parts of the body. Access controls imposed at the level of the file system can therefore only permit or deny access to the entire database, not relevant portions of the data within it. The situation is further complicated by the co-relational nature of this data—many data points have dependencies that prevent them from being processed independently (Pradhan et al., 2007). For example, the positions of the left tibia and left foot at any given time are non-independent because they are physically connected. This means that least-privilege privacy protections for this data require a form of *inference control* (Thuraisingham et al., 1993) in order to be effective and practical. Inference controls allow classified data to be consulted during query processing, but in such a way that the disclosed results do not reveal the values of the confidential data.

To address these challenges, we take the approach of restructuring 3D motion capture data as a relational database equipped with an inference control system. The structure of our database is inspired by past work that has applied indexing structures and hierarchical approaches to reorganize motion capture data in an effort to reduce query processing times and storage overheads for large datasets (Pradhan et al., 2007; Li et al., 2004). In the case of human motion captures, the data are divided into 5 main hierarchies—*viz.* pelvis, right hand, left hand, right leg, and left leg. We leverage this structure to efficiently store the data in a relational database. The database also includes patient identity information, such as name and social security number, and other textual medical information, such as the name of the doctor who attended the patient, the patient's medical history, and treatments given.

The inference controller for the system allows attributes within the database to be classified at various different confidentiality levels. For example, pelvis motion data receives a higher classification than left leg motion because much of the motion data for the whole body can be inferred from the pelvis motion, whereas comparatively little data of the other hierarchies can be inferred from left leg motion. Similarly, patient identity information, such as name and social security number, receive a high classification.

The remainder of the paper proceeds as follows. We begin with a brief overview of related works in Section 2. Section 3 presents our relational database model. Inference problems for this model are discussed in Section 4, and an inference controller that addresses these issues is presented in Section 5. Finally, future work and conclusions are discussed in Section 6.

2 RELATED WORK

Human motion analysis is a longstanding, highly active field of computer vision research (Moeslund et al., 2006). While a significant portion of this research regards physical security problems such as machine-assisted surveillance (Hu et al., 2004), no prior work to our knowledge has considered the problem of enforcing fine-grained access and inference controls over captured motion data.

Relational models for data storage and retrieval have remained a well-established paradigm for organizing large databases for at least the past forty years (Codd, 1970). These organize data into tables that expose important relations between the data points; for example, all data related to a particular patient might be grouped according to the patient's last name.

Aside from facilitating easy manipulation of the data, relational models facilitate more precise enforcement of data security (Staddon, 2003). For example, inference controllers for relational databases leverage the data relations exposed by the model to prevent unauthorized users from inferring private data from public data (Thuraisingham et al., 1993). *Query modification* has been used in the past to enforce discretionary inference and access controls (Stonebraker and Wong, 1974), and later to enforce mandatory controls (Dwyer et al., 1987). The approach enforces policies by extending queries with security constraints prior to query processing. This modular implementation strategy has the advantage of allowing the enforcement mechanism to be developed largely separately from the underlying database management system, affording a convenient separation of concerns.

Our work is inspired by recent advances in human motion data storage that have achieved faster query responsiveness by organizing the data into hierarchical structures that mirror the physical structure of the human body (Pradhan et al., 2007). Organizing the data according to these physical relationships exposes natural co-relations in the data that can be used to more efficiently answer queries involving sub-body motions (e.g., identifying similar knee motions) as well as whole-body motions. We observe that such reorganization leads to natural representation of the data within a relational database, facilitating the implementation of fine-grained inference and access controls over the data through automated query modification.

3 A RELATIONAL DATABASE FOR 3D MOTION CAPTURE FILES

As depicted in Figure 1, human motions can be divided into five sub-body motions—pelvis motions, right hand motions, left hand motions, right leg motions and left leg motions—with the pelvis motions at the root of the tree (Pradhan et al., 2007).

Sub-body motions within each of these classes can be further indexed into additional sub-parts. For example, femur, tibia, foot and toe are sub-parts of the leg. As the figure shows, both left leg and right leg are joined to the pelvis. In the same way, clavicle, humerus, radius, and hand are sub-parts of the arm. Thorax and head are also in the hierarchy, where left arm and right arms are joined to the thorax.

The data hierarchy in Figure 1 mirrors physical relationships in the human segment structure depicted in Figure 2.

To organize 3D motion capture data according to the hierarchy in Figure 1, we partition the data into five separate sub-body motion files (*viz.*, pelvis, right arm, left arm, right leg, and left leg). The location and name of each file can be expressed as a relation in a relational database, as shown in Table 1. The social security number of the patient to whom the motion data belongs is used as the primary key. The patient's and attending physician's names are also included. Any other information about the patient and the doctor are stored in other relations through a standard normalization process.

4 INFERENCE PROBLEMS

The relational database described in Section 3 contains a mixture of various types of data that range over various privacy levels with respect to different principals. For example, a patient's general clinician may need to have access to all of the patient's medical data, but perhaps not the patient's financial or health insurance data. Specialists to whom the general clinician refers the patient may only need access to a subset of the medical data. In addition, privacy violations are possible even when unauthorized principals do not have direct access to confidential data. To illustrate, we consider three representative scenarios that motivate the need for an inference control system for the database. In what follows, we assume the existence of an appropriate authorization mechanism (e.g., based on passwords or keycards) that assigns electronic credentials to users.

Figure 1: A hierarchical tree structure of human body segments

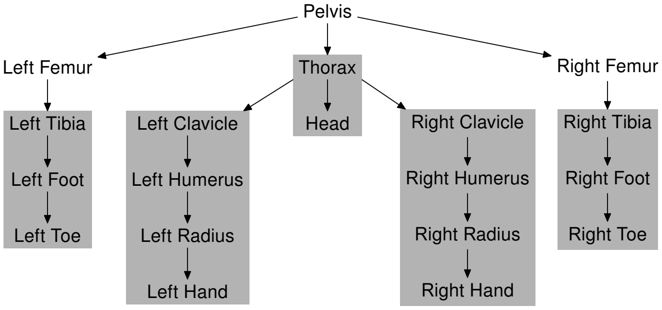

Figure 2: Segment structure for human body with five major sub-body parts

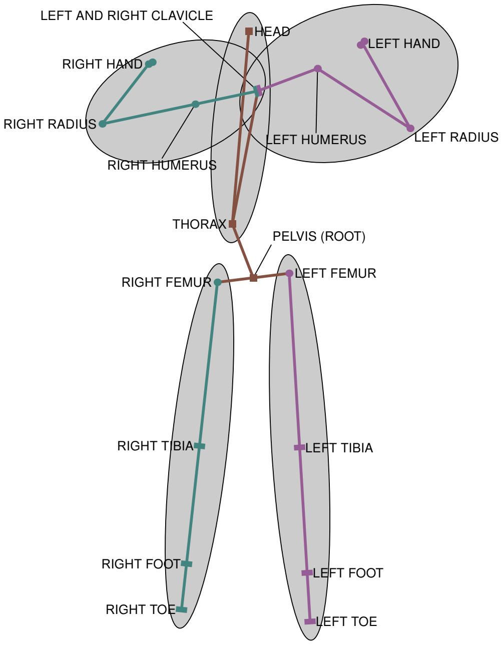

Table 1: Storage of patient motion data within a relational database

SSN	**Patient Name**	**Pelvis File**	**Right Arm File**	**Left Arm File**	**Right Leg File**	**Left Leg File**	**Doctor Name**
1111	John	1011.csv	1012.csv	1013.csv	1014.csv	1015.csv	David
2222	Paul	1021.csv	1022.csv	1023.csv	1024.csv	1025.csv	Michael

Scenario 1

With respect to the 3D motion files, a patient's whole-body motions should be viewable only by the general clinician who treated the patient. Other doctors (e.g., specialists) may only view sub-parts of the patient's motion file. However, if a specialist may use multiple independent queries to separately retrieve all parts of the motion data, he can join them to infer the whole-body motion file. This is known as inference by *semantic association*. Inference controls must therefore generally consider not only the current query but the history of past queries submitted by any given user in order to detect privacy violations.

Scenario 2

Personal, non-medical information about a patient, such as the patient's address and payment plan, are accessible by billing services personnel but not doctors. However, if a doctor is permitted access to the patient's social security number, that information can be used to obtain much of the patient's private information through external sources. This is known as inference by *heuristic reasoning*. Thus, such information must receive a high confidentiality classification even when it does not directly include private information.

Scenario 3

Patients with similar motion data may receive similar diagnoses. Likewise, patients with similar diagnoses may exhibit similar motions. Thus, access to motion or diagnosis data may permit full or partial inference of the other. Similarly, personal characteristics such as height, weight, and gender are often inferable from motion data. Such inferences are known as *analogical reasoning*. To protect patient privacy, attributes that are correlated in this way must receive similar classification levels.

5 THE DATABASE INFERENCE CONTROLLER

Our database inference controller allows a system administrator to protect against the privacy violations described in the previous section by specifying a data privacy policy in the form of privacy constraints. Each constraint assigns a classification level to a multiset of data attributes. Letting A be the set of all attributes and C be the set of all classification levels, a constraint set can therefore be formalized as a relation $\rhd$ of type

$$\rhd : C \times A^{\mathbb{N}}$$

where $A^{\mathbb{N}}$ denotes the space of all multisets over domain A. For example, the relation

$$\mathit{secret} \rhd [\mathit{SSN}^2, \mathit{pelvis}]$$

asserts that *secret* clearance is required in order to access two social security numbers and a pelvis motion file over the full history of database transactions.

The inference controller enforces a policy that requires each principal p to have a clearance level adequate to access every submultiset S of its multiset $H(p)$ of historical past accesses. Formally, principal p must have at least minimal clearance level $MC(H(p))$ defined by

$$MC(A) = \bigoplus_{C \rhd S \subseteq A} C$$

where operator $\oplus$ denotes the *join* of the various required classification levels in the lattice of security labels (Sandhu, 1993). (Intuitively, the join of a set of classification levels is the lowest level that equals or supersedes them all.) If an impending access would expand multiset $H(p)$ so as to violate this requirement, the access is denied.

The design of the inference controller consists of two parts: a *query modifier* and a *response processor*. The query modifier extends the query entered by the user with additional privacy constraints. This modified query is then given to the database for processing. The response processor tracks query responses so as to record the history $H(p)$ of attributes released to each principal. If the new history violates the privacy policy defined by relation $\rhd$ (formalized above), the response is censored and the history remains unchanged.

Privacy Constraints

To test our inference controller, we enforced a policy consisting of four kinds of constraints, summarized below.

Simple constraints place high classifications on attributes that directly contain private data, or that can be used to access private data using external sources, as described in Scenario 2 of Section 4. In our relation we specify three simple constraints:

(i) $secret \rhd [SSN]$,

(ii) $secret \rhd [patient\_name]$, and

(iii) $secret \rhd [doctor\_name]$.

These specify that principals must have secret clearance in order to access any patient social security numbers, names, or attending doctors' names, respectively.

Content-based constraints classify attributes based on co-relations that could be used to infer private data. In our motion capture database, pelvis motion data can be used to infer much of the other motion data, so receives a high classification:

(iv) $secret \rhd [pelvis]$

This constraint was selected because of the position of the pelvis file at the root of the hierarchy tree in Figure 1. Thus, the contents of the pelvis file are considered to be more private than those of any other sub-parts. This addresses Scenario 1 of Section 4.

Association-based constraints classify associations between attributes, and therefore relate non-singleton multisets to classification levels. We have two such constraints in our relation:

(v) $top\_secret \rhd [pelvis, right\_arm, left\_arm, right\_leg, left\_leg]$, and

(vi) $top\_secret \rhd [patient\_name, doctor\_name]$

where *top_secret* is a level above *secret* in the security lattice. This information revealed together is considered a patient privacy violation.

Aggregate constraints limit accesses based on quantity. Such constraints relate classification levels to multisets containing multiplicities greater than 1. We define one aggregate constraint:

(vii) $top\_secret \rhd [head^4]$

This constraint asserts that acquiring 4 head motion files requires a clearance level of *top_secret* or above.

Query Modifier

Our query modifier filters the attribute lists in the SELECT and WHERE clauses of queries so as to satisfy the security policy with respect to the user's clearance level C. That is, given a query from user p, the query modifier replaces attribute list A with a maximal subset $B \subseteq A$ satisfying $MC(H(p) \cup B) \leq C$. If the original query is policy-satisfying, then subset B is simply the original set A, and the query remains unchanged. In the worst case, subset B is reduced to the empty set, and the query is thereby rejected entirely. The following examples illustrate the process.

Example #1: SELECT *SSN* FROM *dbase*
This query is rejected by the query modifier for users not possessing *secret* clearance, and accepted without modification for those that do.

Example #2: SELECT * FROM *dbase*
The * wildcard is first expanded into the full list of attributes. If the user lacks secret clearance, attributes *SSN*, *patient_name*, and *doctor_name* are subsequently dropped due to simple constraints i–iii in Section 5. Next, the *pelvis* attribute is removed due to content-based constraint iv. Finally, history $H(p)$ is consulted to determine whether user p has already reached the limit on head motion files imposed by aggregate constraint vii. If so, *head* is also dropped from the attribute list.

Alternatively, if the user possesses *secret* but not *top_secret* clearance, constraints i–iv do not apply, but *pelvis* and *patient_name* are both dropped in order to satisfy association-based constraints v and vi. Finally, aggregate constraint vii is applied as above.

Example #3: SELECT *pelvis* FROM *dbase* WHERE *patient_name*='John'
If the user p has *secret* clearance for patient John, history $H(p)$ is consulted to determine whether p has already accessed *right_arm*, *left_arm*, *right_leg*, and *left_leg* (which would violate constraint v), or has already accessed *doctor_name* (which would violate constraint vi). If so, the query is rejected; otherwise it is permitted unchanged.

Response Processor

The response processor filters query results based on aggregate constraints, and tracks histories $H(p)$. For instance, the query in example #2 above could potentially yield more than 4 head files. If so, the response processor filters one or more of them in order to bring the result into compliance with aggregate constraint vii.

The filtered query results are then used to update history multiset $H(p)$. In general, history multisets are tracked as a *multiplicity function* m of type

$$m : (P \times A) \to \mathbb{N}$$

where P is the universe of principals and A is the set of all attributes. Maintaining such a function therefore requires $O(|P|\,|A|)$ space.

6 CONCLUSION AND FUTURE WORK

Traditional database inference controllers are designed for multilevel-secure database management systems (MLS/DBMS), but 3D motion captures are typically stored as flat, monolithic files, frustrating the application of traditional inference controls. Our work proposed and tested a strategy for restructuring 3D motion capture data so as to facilitate effective inference control within a standard relational database. Our solution brings together two heretofore disconnected bodies of prior research—relational database security and hierarchical indexing schemes for human motion data. Our choice of query rewriting and response processing as an enforcement strategy allowed us to implement our prototype atop an existing MLS/DBMS system. This avoided the potentially time-consuming and expensive undertaking of redesigning a large-scale MLS/DBMS system to support motion capture data directly.

Our prototype motion capture relation encoded only the top hierarchical level of the index tree described in Section 3. However, past work has investigated indexing schemes with up to 19 different partitions corresponding to various body parts (Pradhan et al., 2007; Li et al., 2004). Applying this more detailed partitioning would open opportunities for more fine-grained policy constraints that distinguish and relate each of these sub-parts. This is an option we intend to explore in future work.

Our work has focused on inference control based on access control. One critical aspect of motion capture data processing is ensuring the privacy of the individuals whose motions are being captured. For example, an adversary should not be able to exploit access to limited motion data to infer diseases that the victim may have, such as arthritis. In addition, motion data collected in a private, medical context should not be divulged to insurance companies without the patient's consent. An important category of future work therefore involves careful formulation of real policies necessary to ensure the privacy of individuals.

REFERENCES

Campbell, T. E., Munro, B. J., Wallace, G. G., and Steele, J. R. (2007). Can fabric sensors monitor breast motion? *Journal of Biomechanics*, 40(3):3056–3059.

Codd, E. F. (1970). A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387.

Dwyer, P. A., Jelatis, G. D., and Thuraisingham, B. M. (1987). Multi-level security in database management systems. *Computers & Security*, 6:252–260.

Fleisig, G. S., Andrews, J. R., Dillman, C. J., and Escamilla, R. F. (1995). Kinetics of baseball pitching with implications about injury mechanisms. *American Journal of Sports Medicine*, 23(2):233–239.

Hu, W., Tan, T., Wang, L., and Maybank, S. (2004). A survey on visual surveillance of object motion and behaviors. *Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 34(3):334–352.

Kaneoka, K., Koshiro, O., Inami, S., and Kayashi, K. (1999). Motion analysis of cervical vertebrae during whiplash loading. *Spine*, 24(8):763–769.

Li, C., Pradhan, G., and Prabhakaran, B. (2004). Indexing of variable length multi-attribute motion data. In *Proceedings of the 2nd ACM International Workshop on Multimedia Databases (MMDB)*, pages 75–84.

Lucareli, P. R., Lima, M. O., Lima, F. P. S., de Almeida, J. G., Brech, G. C., and D'Andréa Greve, J. M. (2011). Gait analysis following treadmill training with body weight support versus conventional physical therapy: A prospective randomized controlled single blind study. *Spinal Cord*, pages 1–7.

Moeslund, T. B., Hilton, A., and Krüger, V. (2006). A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding — Special issue on modeling people: Vision-based understanding of a person's shape, appearance, movement, and behavior*, 104(2–3):90–126.

Morris, M. E., Huxham, F., McGinley, J., Dodd, K., and Iansek, R. (2001). The biomechanics and motor control of gait in Parkinson disease. *Clinical Biomechanics*, 16(6):459–470.

Pradhan, G. N., Li, C., and Prabhakaran, B. (2007). Hierarchical indexing structure for 3D human motions. In *Proceedings of the 13th International Multimedia Modeling Conference (MMM)*, pages 386–396.

Rietman, J. S., Postema, K., and Geertzen, J. H. (2002). Gait analysis in prosthetics: Opinions, ideas and conclusions. *Prosthetics and Orthotics International*, 26(1):50–57.

Saltzer, J. H. (1974). Protection and the control of information sharing in Multics. *Communications of the ACM*, 17(7):388–402.

Sandhu, R. S. (1993). Lattice-based access control models. *IEEE Computer*, 26(11):9–19.

Staddon, J. (2003). Dynamic inference control. In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD)*, pages 94–100.

Stonebraker, M. and Wong, E. (1974). Access control in a relational data base management system by query modification. In *Proceedings of the 1974 ACM Annual Conference*, pages 180–186.

Thuraisingham, B., Ford, W., Collins, M., and O'Keeffe, J. (1993). Design and implementation of a database inference controller. *Data & Knowledge Engineering*, 11(3):271–297.