The Heat Based Rump Descriptor for Identification of Very Thin Goats in Dairy Farms

Susana Brandão[1], Ana Vieira[2], João P. Costeira[3], Manuela Veloso[4]

***Abstract*— In this paper we address the problem of identifying goats in dairy farms with sub-optimal values of Body Condition Score (BCS). The BCS conveys information on whether an animal is fat or thin and its identification has a strong economic impact as very thin animals have poorer milk production and associated health problems. Albeit its important implications, not only there is no automated way of assessing the BCS in dairy farms, but current available techniques require specially trained personnel. However, the recently introduced Pictorial Scale for BCS assessment in dairy goats shows that the rump region has several visual cues strongly correlated with animal's BCS values. In this paper we move towards the automatic assessment of BCS by developing a descriptor for rump's 3D surfaces, collected by an RGB-D camera. The use of 3D surfaces as the basis for identification is fundamental, as it allows data collection without requiring animal handling to ensure careful alignment between camera and animal. However, the identification of the rump region in the 3D surfaces is very difficult, which leads to a large variability in the type of surfaces associated with the same BCS value. The descriptor we here introduce, the Heat Based Rump Descriptor (HBRD), uses diffusion geometry concepts to seamlessly handle the difficulty in defining a rump region and the resulting large variability of shapes. We test our descriptor in a dataset of 32 dairy goats and show that our descriptor is able to effectively cluster all the very thin animals.**

I. INTRODUCTION

The Body Condition Score (BCS) is correlated with an animal fat deposits and is an important animal-based indicator of animal welfare. Furthermore, very low BCS, as those represented in Fig. 1(a), are also correlated with low milk production [1] and are not in adherence with consumers expectations on animal's welfare [2].

European Union, having recognized the farm animals' right of freedom from hunger and thirst, is currently moving towards the introduction of BCS as a key indicator on welfare assessment protocols on goat farms. However, standard techniques for estimating the BCS in goats , e.g., as those presented in [8], cannot be used in large scale assessments, as they require restraining and handling of each animal individually by specially trained assessors.

The recently introduced Visual Body Condition Scoring System, [15], addresses the scalability problem by creating illustrations to allow non-experts to assess the BCS by visual inspection. For the construction of the Pictorial Scale, authors identified several visual features in the rump region that are strongly correlated with the animal BCS. Those features correspond to distances between bones and muscle folds, which are easy to visually identify. The features were used to define a *standard* individual of each class, from which a professional illustrator generated drawings for the Pictorial scale. The Pictorial scale can now be used in farms, but still requires trained evaluators.

(a) Very thin (b) Normal (c) Very Fat

Fig. 1. Examples of very thin, normal and very fat animals at different distances and orientations from a RGB-D camera. Each animal was manually evaluated to assess its BCS score.

The features identified in [15] worked well for the purpose of creating visually accurate illustrations. However, to retrieve such features, authors acquired photographs taking careful control on conditions such as: i) animals' stillness; and ii) rumps' alignment with the camera, which are difficult to ensure without animal handling. In this paper we move towards a scenario where no handling is required by using RGB-D cameras, as 3D information handles better changes in the orientation between camera and animal. Such cameras can be fixed on top of the animals' normal path, and can accurately collect data at roughly 2m from the animal.

RGB-D cameras, such as the Kinect camera, provide both an RGB image and a depth image, from which we can recover 3D surfaces corresponding to the animal surface. From the whole animal, we extract the rump as showed in Fig. 2 using a manual labeling approach similar to the one presented in [15].

As noted in [15], the main difference between different

[1]Susana Brandão is with ECE at CMU, USA and IST-UL, Portugal. sbrandao@ece.cmu.edu

[2]Ana Vieira is with Veterinary School at FMV-UL, Portugal. ana.lopesvieira@gmail.com

[3]João P. Costeira is with Faculty with ECE at IST, Portugal. jpc@isr.ist.utl.pt

[4]Manuela Veloso is with Faculty with CS, CMU USA. mmv@cs.cmu.edu

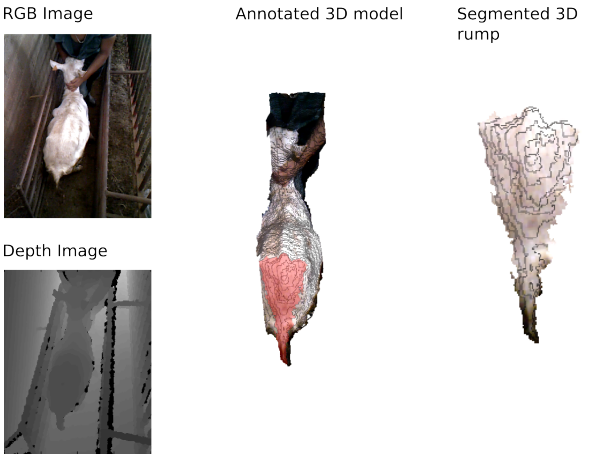

Fig. 2. Acquiring rumps' 3D surfaces.

BCS categories are the fat reserves in the rump, which yield a bulkier appearance in fatter animals. To correctly assess the animal class, we focus on descriptors that represent changes in the volume between rumps of different animals. Furthermore, the most noticeable changes in the rump volume concern its upper part, near the hip.

However, the direct comparison of volume between rumps 3D surfaces is very challenging, as: (i) rump shapes vary considerably between animals, regardless of BCS, as showed in Fig.1; and (ii) it is difficult to consistently define the rump region in a meaningful and consistent way.

Without a clear definition of the region of interest, we cannot compare two surfaces. In fact, from the several descriptors available for the recognition of 3D shapes,[16], [10], [3], none that we are aware of handles the problem of representing a shape that itself is ill defined. We thus propose to assess the changes in volume from one surface to another based on how much they differ from a plane. Animals with rumps that are more similar to a plane, have smaller fat deposits, i.e., are thinner. By introducing an intermediary shape, we avoid mapping and registration between rumps.

To capture changes in this vague region of the upper part of the rump, without having to specifically segment it, we use multi-scale descriptors, i.e., descriptors that provide information on how a given point in the surface is related to the whole surface by considering increasingly large neighborhoods for that point. An example of such descriptors are Heat Kernel Signatures [13].

Heat Kernel Signatures and other heat based descriptors describe how connected a point is to its neighborhood by simulating heat propagating over a surface. In Fig. 3 we show several snapshots of the process of heat propagating from an initial heat source to the whole surface. At each fixed time instant, the temperature of a point in the surface is related to its distance to the initial source. Furthermore, as time passes, the temperature in points further from the source increases, while the temperature at the source decreases. The change in the temperature is more significant in the first instants, when there are sharp contrasts over the surface, than in the end, when the temperature over the whole surface becomes constant, regardless of the surface shape.

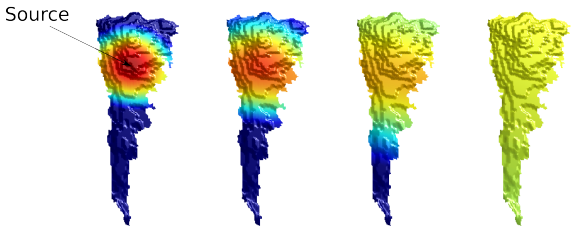

Fig. 3. Heat diffusion over a goat's 3D rump. Red corresponds to higher temperatures and blue to colder ones.

Thus, heat diffusion has two characteristics that make it the ideal choice of representation: (i) naturally introduces a notion of scale, and (ii) temperatures can be used as a surrogate to distances[3], specially when surfaces are noisy and have a poor resolution such as those from common RGB-D cameras.

In this paper we move towards the automatic identification of the body condition score of farm goats by introducing a Heat Based Rump Descriptor (HBRD) that:

- represents regions at different scales, allowing to focus on the upper part of the rump, without having to explicitly segment the region;
- describes the rump by comparing it against a default shape, in this case a plane.

Such descriptor allows to handle the variability in the animals' shape and the difficulty in defining the region of interest.

In the following we show how we obtain an initial segmentation of the rump region given the output of an RGB-D sensor, and we provide full detail on how to estimate HBRD descriptors in any given rump.

II. Data Acquisition

While leaving the milking room, animals pass one by one through a narrow corridor. We placed a calibrated RGB-D sensor on a fixed point above the animals' path. An expert manually evaluated the animals' BCS to provide ground truth using the simplified 3 points scale defined in [15].

While we cannot identify accurately the rump region in the different animals, we follow [15] and define the region based on the rump bone structure, namely the tuber sacrale (hip or hook bones) and the tuber ischia (pin bones), represented in Fig. 4(a). The tips of these bones correspond to easily identifiable features in the RGB images of animals of all categories, as we show in Fig. 4(b)-(d).

From the depth image, D, we can recreate the goat 3D surface, as illustrated in Fig. 2. The surface corresponds to a set of triangles, represented by a list of vertex coordinates $X = [\bar{x}_1^T, \bar{x}_2^T, ..., \bar{x}_N^T]$ and a set of edges $E = \{e_1 = (1,2), ..., e_{N_e} = (k,N)\}$. Each vertex corresponds to a pixel in the depth image, and the coordinates are obtained by calibrating the sensor. We construct the set of edges based

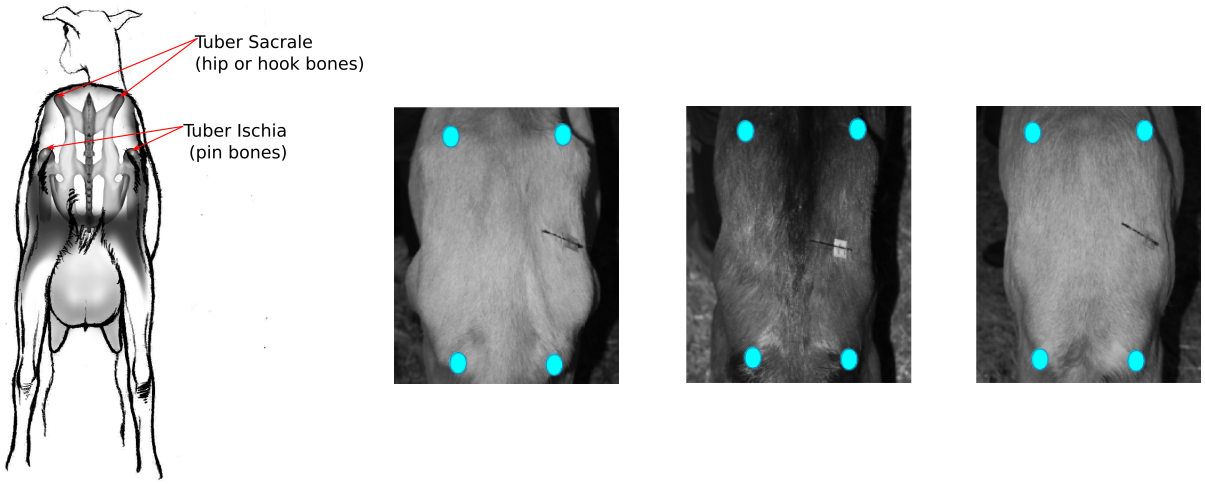

(a) Detail on the rump bone structure. (b) Examples of photographs with annotated hip and pin bones.

Fig. 4. Rump identification: Detail on the bone structure of a goat rump, showing that hip and pin bones are part of the same structure and lay on the same plane, and photographs showing that the bones are easy to identify.

on adjacency relations between pixels in the depth image, creating a mesh of triangles that cover the surface, without overlap. From the camera calibration, we can also map the annotations in the RGB image, I, to the depth image, D, to obtain the 3D coordinates of the left and right hip bones, $\bar{b}_{l,r}$, and pin bones $\bar{p}_{l,r}$.

When the goat is standing, the four bone tips approximately define a plane, as the hip and pin bones are rigidly connected. By finding the orientation of the plane defined by the four bone tips with respect to the floor, we rotate the whole surface, so the bone tips lay in the $x-y$ plane. We define the rump as all the points with a positive z. This segmentation is reproducible and consistent, albeit it may lead to the inclusion of other parts of the animal in the rump, e.g., the tail.

To account for changes in the animal size, we normalize both x and y coordinates of all vertices, so that the bone tips of all the animals are in the same position $\bar{h}'_{l,r}$ $\bar{p}'_{l,r}$ in the $x-y$ plane. To account for possible hip or tip bones miss-alignment, we normalize using a projective transformation. The resulting normalized coordinates, $X_{norm} = [\bar{x}_1^{norm} = [x_1^{norm}, y_1^{norm}, z_1], ..., \bar{x}_N^{norm}]$, maintain the same z-coordinate. The edges in the normalize surface connect the same vertices as the edges in the original one.

After segmentation and normalization, we obtain a set of rumps similar to those represented in Fig. 5.

III. Rump Description

A. Representing variable surfaces

Rumps in Fig. 5 highlight that the most distinct feature among all rumps is that thin goats are almost flat. The rumps also illustrate the intra-class variation resulting from the natural variability of goats shapes and sizes. In particular, it shows that goats have different features that are not related with the BCS, e.g., rump boundaries change considerably across animals, and in some animals the tail is included in our estimation of the rump region. We must also account for errors in the segmentation process, such as (i) there is a large uncertainty in the identification of hip and pin bones on the animals rump, (ii) it is difficult to ensure that the bone tips are in a plane, and (iii) errors in camera calibration result in errors in the map between RGB and depth images.

Common approaches for 3D shape representation, such as bag of features or 3D holistic representations are not effective in describing these variations, as they all assume that any input shape is fully explained by the category they want to represent and eventually classify. As far as we are aware, there is no previous work on the representation of 3D shapes where the shape itself was not explicitly defined.

We compare the differences in volume by extracting shape information, e.g., distances between points and areas, and compare it with the same information extracted from a planar projection, as showed in Fig. 6. The planar projection corresponds to the same mesh, but with z-coordinate set to zero, $X_{plane} = \left[\bar{x}_1^{plane} = [x_i^{plane}, y_i^{plane}, 0], ..., \bar{x}_N^{plane}\right]$.

The comparison between the two surfaces is possible because there is a natural bijection relating the two surfaces, i.e., there is an one to one relation between points in the rump and in the planar projection. We thus compare the two surfaces by computing a geometry dependent function at each point and compare the values of both surfaces at related points.

As stated, in this work we use the temperature resulting from a heat diffusion process, as it provides a natural segmentation of the interest region and depends on the geometry of each surface, as it occurs faster in planar surfaces. Other functions, e.g., the distance to a point, also depend on the geometry of each surface, however require a rigorous definition of the interest region. We thus assess if the geometry of the two surfaces is similar by comparing the

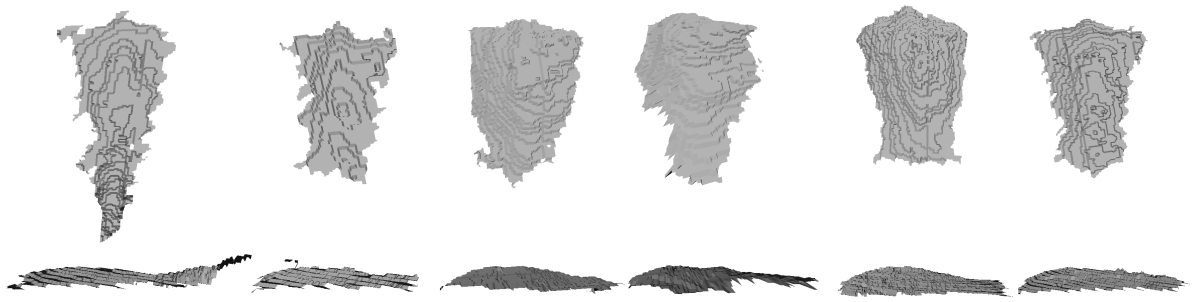

(a) Thin (b) Thin (c) Normal (d) Normal (e) Fat (f) Fat

Fig. 5. Example of rumps from different animals. The top image represent a view from the z-axis, while the bottom view from the x-axis.

Fig. 6. Example of a planar rump, on the left, built from the regular rump on the right.

temperature at equivalent points in both surfaces.

B. Heat Diffusion on Discrete Surfaces

Heat based descriptors have shown good results at representing surfaces retrieved from depth sensor [3], [4], [5] and other 3D shapes [13], [6]. We here briefly review the necessary steps to compute a temperature $\bar{T}(t) \in \mathbb{R}^N$ on all the vertices in the surface, at each time instant t. The familiar reader may skip to section III-C.

Heat diffusion in discrete surfaces, such as the one obtained from depth images, is described by eq.1, [9]

$$\partial_t \bar{T}(t) = -L\bar{T}(t) \tag{1}$$

where L is the discrete Laplacian matrix, which is related with the Laplace-Beltrami operator defined in continuous surfaces [13]. Such operator returns the temperature second derivative as defined over the surface, i.e., taking into account that the surface is not necessarily a plane.

The discrete version we use in this work is associated with a graph interpretation of the organized set of points in the depth image. As showed in Fig. 7, each pixel i in the depth image leads to a vertex in the graph with coordinates $\bar{x}_i$. The vertices are connected by the triangle edges E and to each edge $e = (i, j)$ connecting a vertex i to a vertex j there is an associated weight $w_{i,j} = \|1/\|\bar{x}_i - \bar{x}_j\|^2$.

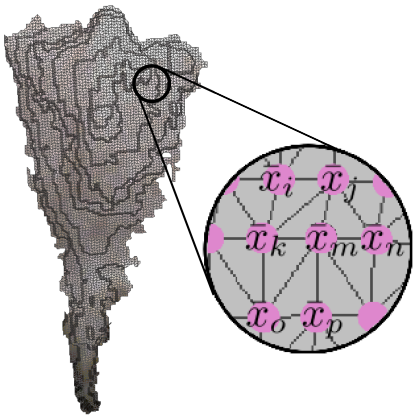

Fig. 7. Graph structure of surfaces retrieved with an RGB-D camera.

For a graph/surface with N vertices, the Laplacian is the $N \times N$ symmetrical matrix $L = D - W$, with $[W]_{i,j} = w_{i,j}$ if there is an edge connecting the vertex i and j, i.e., if $e = (i, j) \in E$, and 0 otherwise. D is a diagonal matrix, with $[D]_{i,i} = \sum_{j=1}^{N} [W]_{i,j}$. The resulting L matrix, using the above definition for the weights, corresponds to the finite differences approximation to the second derivative.

With Newman boundary conditions and an initial temperature over the surface equal to $\bar{T}(0)$, the temperature at any other time instant $\bar{T}(t)$ can be written in close form with respect to the eigenvalues λ_i and eigenvectors $\bar{\phi}_i$ of L:

$$\bar{T}(t) = \sum_{i=1}^{N_V} \bar{\phi}_i \exp\{-\lambda_i t\} \bar{\phi}_i^T \bar{T}(0). \tag{2}$$

C. Heat Based Rump Descriptors

We evaluate how much a rump differs from a plane by considering a heat diffusion process starting at its center and the equivalent vertex on its planar projection. Thus, the initial condition for both the temperature in the normalized surface $\bar{T}(0)$ and in the plane, $\bar{T}'(0)$ will be equal to each other and be zero everywhere except for some vertex c in the center of the rump, i.e., $[\bar{T}]_c = 1$ and $[\bar{T}]_i = 0 \ \forall i \neq c$.

The vertices at the center of both rumps, with coordinates $\bar{x}_c$, and $\bar{x}_{plane,c}$, are those closest to the center of the quadrilateral defined by $\bar{h}'_{l,r}$, $\bar{p}'_{l,r}$ in both X_{norm} and X_{plane} respectively.

For each animal, given the set of edges E and the two sets of vertex coordinates X_{norm} and X_{plane}, we compute the Laplacian for each surface, L_{norm} and L_{plane}. From each Laplacian we compute the first 300 eigenvectors and eigenvalues and, given the initial condition, $\bar{T}(0)$, we propagate the temperature at both surfaces using eq. 2. As to each point in the original surface corresponds a single point in the planar surface, we can compute the difference between the temperature at both surfaces, $\Delta\bar{T}(t) = \bar{T}_{norm}(t) - \bar{T}_{plane}(t)$ at each time instant.

We evaluate the time difference at exponentially large time intervals, as changes in temperature occur faster in the beginning. In particular, we use time instants $t_k = 0.1e^{-k\delta t}$, spanning from $1/700$ to $1/10$. We focus on the rump upper part by assessing $\Delta\bar{T}(t)$ at a subset of vertices $\mathscr{S}$ that form the shortest path in the planar mesh between $\bar{x}_c$ and $\bar{h}'_l$, which we compute using the Dijkstra algorithm [7].

Finally we construct the descriptor, $\bar{z}$ by considering, for each time instant t_k, the maximum of $\Delta\bar{T}(t_k)$ over the subset of vertices $\mathscr{S}$, i.e.,

$$\bar{z} : [\bar{z}]_k = \max_{x \in \mathscr{S}} [\Delta T(t_k)]_x \tag{3}$$

The main steps for computing HBRD are highlighted in Algorithm 1. The algorithm requires as input an RGB image, I, a Depth image, D, which we here assume that is already mapped into the RGB image. As fixed input parameters, the algorithm further requires the time instants at which we compute the temperature, $\bar{t}$, and the coordinates of the left and right hip and pin bones in the normalized rump, $\bar{h}'_{l,r}$, $\bar{p}'_{l,r}$. In this study, the position of the bone tips in the RGB image $\bar{h}_{l,r}$, $\bar{p}_{l,r}$ is provided by the user, however we expect that this step can be automated using feature matching and taking advantage of the 3D information provided by the depth image, as in [12].

Algorithm 1: Heat Based Rump Descriptor (HBRD).

Input: RGB image: I; Depth image: D; Time instants: $\bar{t}$; bone tips in the normalized rump: $\bar{h}'_{l,r}$, $\bar{p}'_{l,r}$

Output: Rump descriptor, $\bar{z}_r$.

Manually Annotate Hip and Pin Bones in the RGB Image:
$[\bar{h}_{l,r}, \bar{p}_{l,r}] \leftarrow \text{annotate}(I)$
Segment and Normalize depth image:
$[X_{norm}, E] \leftarrow \text{segmentNormalize}(D, \bar{h}_{l,r}, \bar{p}_{l,r}, \bar{h}'_{l,r}, \bar{p}'_{l,r})$
$X_{plane} \leftarrow \text{project}(X_{norm})$
Find Path Between Center and Left hip bone:
$\bar{x}_c \leftarrow \text{centroid}(\bar{h}'_l, \bar{h}'_r, \bar{p}'_l, \bar{p}'_r)$
$\mathscr{S} \leftarrow \text{dijkstraShortestPath}(mesh, X_{plane}, \bar{h}'_l, \bar{x}_c)$
for $i = 1; i < size(\bar{t}); i{+}{+}$ **do**
  Estimate both temperatures distributions, from eq 2:
  $\bar{T}^{\mathscr{S}}_{norm} \leftarrow \text{propagateHeat}(X_{norm}, E, \mathscr{S}, [\bar{t}]_i)$
  $\bar{T}^{\mathscr{S}}_{plane} \leftarrow \text{propagateHeat}(X_{plane}, E, \mathscr{S}, [\bar{t}]_i)$
  $\Delta T([\bar{t}]_i) = \bar{T}_{norm} - \bar{T}_{plane}$
  Get descriptor, from eq 3:
  $[\bar{z}_r]_i \leftarrow \max(\Delta T([\bar{t}]_i))$
end

IV. Results

We used Algorithm 1 to describe different animals.

Fig. 8 shows that temperature in thinner animals converges faster to that the planar rump. The figure represents four rumps, two very thin and two normal. The colors represent the absolute difference between the temperature in the rump to the planar rump. The black line in the upper left part of each rump corresponds to the shortest path $\mathscr{S}$.

Fig. 9 shows the descriptors for the animals in Fig. 8. There is a clear difference over the maximum of the difference between normal and thin animals. Furthermore, we note that by looking only into what happens on the top part of the rump, the animal's tail has little impact on the temperature on the top part of the rump.

Finally, we show that HBRD differentiates thin animals among a dataset of 32 animals, 9 thin, 17 normal and 6 fat. Fig. 10 shows the 3D-Isomap projection[14] of the set of descriptors. The Isomap projection, similarly to PCA projections[11] allows to visualize data of high dimension. However, Isomap provides a representation which minimizes distortion of the distances over nonlinear subspaces, in contrast to PCA that assumes linear subspaces and euclidian norms.

Results show that very thin animals are well clustered, i.e., that the Heat Based Rump Descriptor captures a very elusive characteristic. Some supervised machine learning algorithm can then be trained using these descriptors and used used for automatic classification. Support vector machines would be good candidates for classification. We further note that, by introducing a comparison surface that shares with the original rump most of the characteristics that are not intrinsic to the class, e.g., the tail, we obtained a descriptor that correctly represents the class dependent characteristics.

V. Conclusions and Future Work

In this paper we introduced the Heat Based Rump Descriptor (HBRD) for the identification of very thin goats in dairy farms. The identification of such animals is of utmost relevance not only by the economic implications of the decrease in the milk production associated with a low BCS, as it is in direct violation of the animal's rights.

The HBRD assesses the BCS by the rump volume. To handle the large variability of animal shapes and the difficulty of exactly setting the limits of the relevant part of the rump, HBRD uses heat diffusion to represent distances between points in two equivalent surfaces. The volume is assessed by having the surfaces differ only on the characteristic that we want to measure, i.e., the volume. The use of heat diffusion allows to soft segment the region of interest, as the difference in temperature in both surfaces will be more significant in

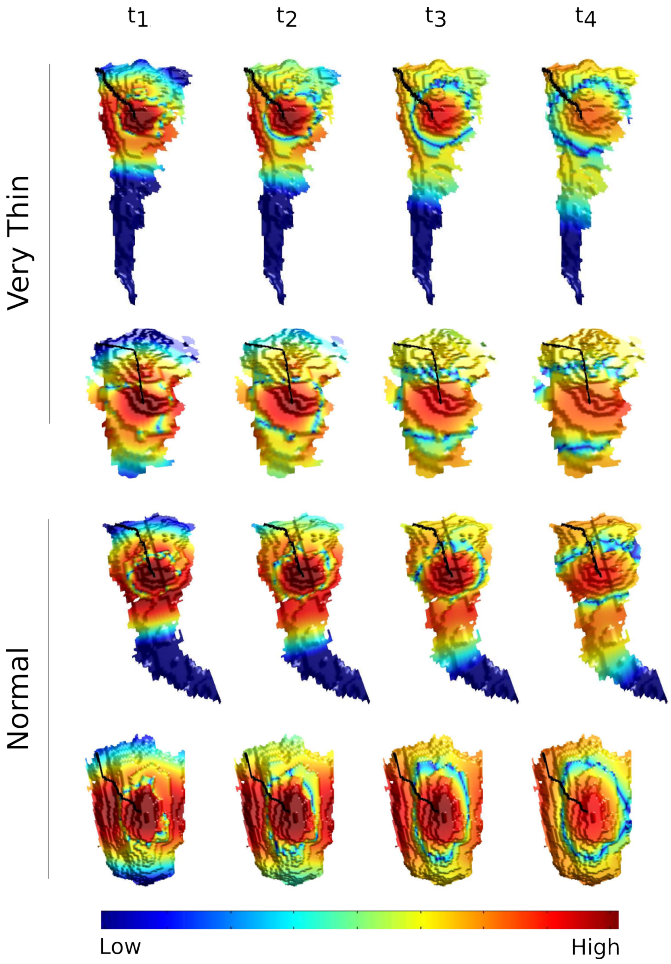

Fig. 8. Difference over time between the temperature over the rump and over the planar rump.

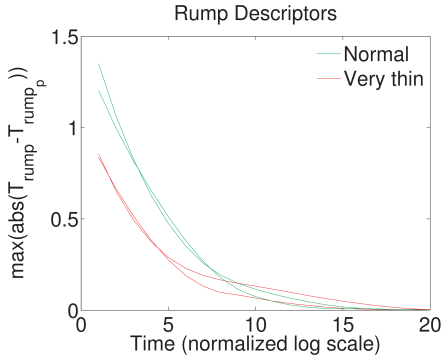

Fig. 9. Maximum difference over time and over the path marked in Fig. 8.

initial time instants, where only the regions close to the source have a significant impact on the temperature.

Using a dataset of 32 animals, we showed that HBRD provides a good representation for the problem, as all the very thin animals in the dataset were clustered together.

By the introduction of relevant descriptors, the work here presented is an important step towards the automation of BCS assessment in dairy goats. Future work should then focus on the automatic identification of the hip and pin bones in the RGB images.

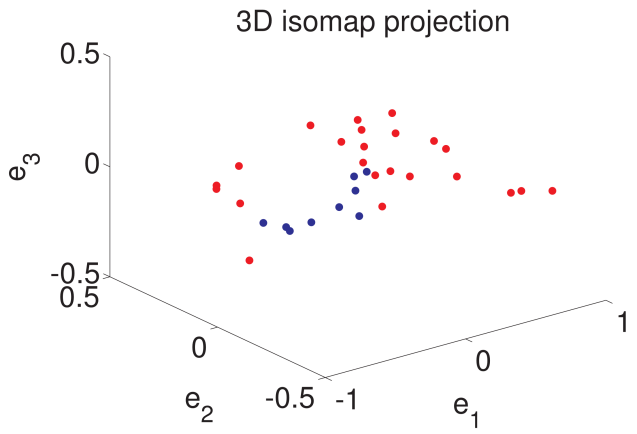

Fig. 10. 3D Isomap projection of the rump descriptors on a dataset of 32 animals. The blue points correspond to thin animals, while red correspond to normal and very fat.

ACKNOWLEDGMENT

The authors would like to thank Telmo Nunes, Inês Ajuda and George Stiwell, from the School of Veterinary Medicine, University of Lisbon for valuable comments and for making possible the data collection for this paper.

REFERENCES

[1] Invited review: Body condition score and its association with dairy cow productivity, health, and welfare. *J. Dairy Sci.*, 2009.
[2] H. Blokhuis, M. Miele, I. Veissier, and B. Jones. *Improving farm animal welfare - Science and society working together: the Welfare Quality approach.* 2013.
[3] S. Brandão, J. Costeira, and M. Veloso. Partial view heat kernel descriptor. In *ICRA*, 2014.
[4] S. Brandão, M. Veloso, and J. P. Costeira. Multiple hypothesis for object class disambiguation from multiple observations. In *3DV*, 2014.
[5] S. Brandão, M. Veloso, and J. P. Costeira. Global localization by soft object recognition from 3d partial views. In *Iros*, 2015.
[6] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *CVPR*, 2010.
[7] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1959.
[8] J. Hervieu and P. Morand-Fehr. Comment noter létat corporel des chèvres. 1999.
[9] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2006.
[10] R. B. Rusu, G. Bradski, R. Thibaux, and J. Hsu. Fast 3D Recognition and Pose Using the Viewpoint Feature Histogram. In *IROS*, October 18-22 2010.
[11] J. Shlens. A tutorial on principal component analysis. In *Systems Neurobiology Laboratory, Salk Institute for Biological Studies*, 2005.
[12] X. Song, J. Schuttea, P. van der Tol, F. van Halsema, and P. Groot Koerkampa. Body measurements of dairy calf using a 3-d camera in an automatic feeding system. In *International Conference of Agricultural Engineering*, 2014.
[13] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *SGP*, 2009.
[14] J. Tenenbaum, V. de Silva, and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 2000.
[15] A. Vieira, S. B. ao, A. Monteiro, I. Ajuda, and G. Stilwell. Development and validation of a visual body condition scoring system for dairy goats with picture-based training. *J. of Dairy Sci.*, 2015.
[16] W. Wohlkinger and M. Vincze. Ensemble of shape functions for 3d object classification. In *ROBIO*, 2011.