an exhaustive search. Therefore, it seems that the improved rekey message format proposed by Trappe *et al.* [2] is actually not as secure as the traditional one (1), and in fact, their two rekey schemes in themselves may be highly vulnerable.

Finally, we note that it is possible to fix the pinpointed weakness with a remedial approach, which is in fact borrowed from the "secure lock" scheme proposed in [7]. To this end, we need to combine the desired *security level* provided by the traditional but partitioned solution depicted in (1) with a favorable *homogenized format* as of the rekey messages depicted in (2) and (4). Both merits can be preserved by employing the mathematics of the Chinese Remainder Theorem (CRT) [7] with the GC constructing a rekey message

$$\alpha_\epsilon \equiv E_{K_i}(K_\epsilon) \pmod{N_i}, \quad 1 \leq i \leq n \qquad (6)$$

where N_i are n public integers that are pairwise relatively prime. On receiving (6) from the GC, each $u_i \in \mathcal{G}$ can derive the root KEK from the homogenized rekey message by computing $K_\epsilon = E_{K_i}^{-1}(\alpha_\epsilon \bmod N_i)$. We refer to [7] for more details. In a nutshell, the GC publishes certain numbers that are pairwise relatively prime, computes a common solution for (6) with the CRT either regarding the entire group $\mathcal{G}$ (when n is relatively small) or mapped to a logical key tree (when n is large), and broadcasts such a homogenized α_ϵ. Then each group member u_i can acquire K_ϵ (and then K_s) by employing a modular operation followed by a symmetric decryption. This way the identifiers in (1) are eliminated, while the security breach in (2) or (4) is also avoided.

REFERENCES

[1] C. K. Wong, M. Gouda, and S. S. Lam, "Secure group communications using key graphs," *IEEE/ACM Trans. Netw.*, vol. 8, no. 1, pp. 16–30, Feb. 2000.

[2] W. Trappe, J. Song, R. Poovendran, and K. J. R. Liu, "Key management and distribution for secure multimedia multicast," *IEEE Trans. Multimedia*, vol. 5, no. 4, pp. 544–557, Dec. 2003.

[3] D. Naor, M. Naor, and J. Lotspiech, "Revocation and tracing schemes for stateless receivers," CRYPTO 2001, *Lecture Notes in Computer Science*, vol. 2139, pp. 41–62, 2001.

[4] J. Lotspiech, S. Nusser, and F. Pestoni, "Anonymous trust: Digital rights management using broadcast encryption," *Proc. IEEE*, vol. 92, no. 6, pp. 898–909, Jun. 2004.

[5] M. J. Mihaljević, M. P. C. Fossorier, and H. Imai, "Security evaluation of certain broadcast encryption schemes employing a generalized time-memory-data trade-off," *IEEE Commun. Lett.*, vol. 11, no. 12, pp. 988–990, Dec. 2007.

[6] W. T. Zhu, "General weakness in certain broadcast encryption protocols employing the remainder approach," in *Proc. IEEE Int. Conf. Communications (ICC) 2008*, May 2008, pp. 1620–1624.

[7] G.-H. Chiou and W.-T. Chen, "Secure broadcasting using the secure lock," *IEEE Trans. Softw. Eng.*, vol. 15, no. 8, pp. 929–934, Aug. 1989.

Synchronization of Olfaction-Enhanced Multimedia

Oluwakemi A. Ademoye and Gheorghita Ghinea, *Member, IEEE*

Abstract—This paper presents the results of an experimental study carried out to explore, from an end user perspective, the temporal boundaries within which olfactory data can be used to enhance multimedia applications. Results show the presence of two main synchronization regions, and that olfaction ahead of audiovisual content is more tolerable than olfaction behind content.

I. INTRODUCTION

MULTIMEDIA synchronization is concerned with defining and maintaining the temporal relationships between two or more correlated media objects that are combined, processed and presented together to produce a multimedia system or application. Because multimedia is essentially about using these multiple media objects to communicate information to users, achieving the desired synchronization between them is vital to the success of these systems [1], [4], [13], [15].

Synchronization has been traditionally associated with quality of service (QoS) parameters such as delay, delay jitter and inter-media skew to allocate resources and ensure that the desired multimedia synchronization quality is provided to the user. In recent times, there has however been a paradigm shift in the community with the recognition that technical parameters alone are not sufficient to produce the highest user-perceived multimedia quality. Such algorithms must also incorporate measures of the user-perceived quality experience.

From a synchronization viewpoint, the foremost study carried out in this respect was the one undertaken by Steinmetz [1], [15]. He investigated user-perceived synchronization requirements for lip synchronization and pointer synchronization, as different synchronization requirement results had been reported by different research groups. The purpose of these experiments was to discover the boundaries within which inter-media skew between audio and video data, for the case of lip synchronization, i.e., between audio and video, and between audio and pointer is tolerable to the user. It was shown that while people are sensitive to the correct timing of combined media, users are still able to tolerate multimedia presentations that are not perfectly synchronized to a certain degree.

Whilst one does not expect the inter-media synchronization requirements between olfactory data and other media objects to be as stringent as that required by lip synchronization, smell has its own peculiarities that will indeed impact on its synchronization requirements, and is the focus of this paper.

II. OLFACTORY DATA

Olfactory data, i.e., smell output information, is virtual data, unlike other media streams which have the ability to be stored in some computer data form. Smell, moreover, has a tendency to linger, and its emission usually happens at a slow pace, since it relies on atmospheric breeze to move it through the air. As such, smell is again different from the media objects multimedia applications have become accustomed to, e.g., audio, visual, text, images and animations, whose component

TABLE I
VIDEO CATEGORIES AND SMELLS USED

Smell Category	Burnt	Flowery	Fruity	Foul	Resinous	Spicy
Video Description	Documentary on bush fires in Oklahoma	News broadcast featuring perfume launch	Documentary about rotting fruits	Cookery show on how to make a fruit cocktail	Documentary on Spring allergies & cedar wood	Cookery show on how to make chicken curry
Smell Used	Burning Wood	Wallflower	Strawberry	Rancid Acrid	Cedar Wood	Curry

Fig. 1. Vortex Active from http://www.daleair.com/acatalog/Dispensers.html.

media streams have a transitory nature and do not experience this natural delay. In addition, the occurrence of data loss with olfactory data is also something to be expected, as sensitivity to the presence of a smell tends to decrease with prolonged exposure to it, i.e., olfactory adaptation or habituation [11].

Further compounding the challenge of olfactory data is the fact that odor perception is highly subjective, has a habit of changing, as well as being greatly influenced by age, sex, social and cultural factors, in addition to emotions, memory, experience and input from other sensory modalities [5], [10], [12], [14]. Thus, the same odor may be perceived differently by separate groups of people or individuals—pleasant to some, unpleasant to another, and neither pleasant nor unpleasant to others.

Furthermore, despite the distinctive uniqueness associated with smell that makes it possible to detect its presence in the air, research has shown that it is easier for people to smell something than it is for them to identify what they have smelt, and more often than not people will try to identify a smell by associating it with a smell they are familiar with [2], [3], [9]. For these reasons, and the lack of standard smell classification schemes, it is advised that familiar smells be used for experiments involving the use of olfactory data [9], [12]. In the following section, we examine how some of the characteristics of olfactory data discussed here may impact on olfactory data synchronization.

III. OLFACTORY DATA SYNCHRONIZATION

Although computer generated smell is time-based, with the ability to explicitly specify its start time and duration, the characteristic nature of smell means that its perception may not be consistent with the objective time. Moreover, as a result of the difficulties associated with odor perception, as well as the virtual nature of olfactory data, controlling the parametric values used for measuring synchronization for olfactory data will be quite difficult, as most of these values are more dependent

on external factors than they are on the multimedia application or computing environment for olfactory data. For example, controlling delay for olfactory data is more dependent on external factors such as the time it takes the output device to emit the required scent or the effects of nature.

Olfactory data itself cannot be transmitted via distributed systems and therefore delay jitter is of little or no relevance for olfactory data, as it will only apply to the commands used to trigger its execution. Bit error and packet loss, like delay, for olfactory data are also more dependent on external factors, Mother Nature and the users' sense of smell in this case. However, inter-media skew is something that *must* be controlled for any media object that needs to be synchronized as achieving the desired temporal relations between the combined media objects is the core purpose of synchronization. As such, inter-media skew is the only measure of real significance that can be used and optimized while determining olfactory data synchronization requirements and it is chiefly what we have controlled to measure the user-perceived experience of the synchronization effect of an olfactory-enhanced multimedia application in our experimental study.

IV. EXPERIMENTAL METHODOLOGY

The goal of our experiment was to discover the tolerable inter-media skew synchronization requirements for combining computer generated scent with related audiovisual content.

A. Experimental Materials

We designed a multimedia presentation display program, which displays visual and audio media content from video clips synchronized with olfactory data (Table I). For this experiment, we have used the Vortex Active scent dispensing system by Dale Air, a U.K.-based company well known for providing themed aroma displays. The Vortex Active (shown in Fig. 1) is a personal computer smell dispensing system which uses miniature fans to propel the emitted smells in the right direction. It connects to the computer via a USB port, and we used the USB fan controller application supplied with the device to control the synchronized release of olfactory data during each video playback.

We opted for six smells related to olfactory research, i.e., the Hans Henning "smell prism" classification scheme [5], [9], as we felt it would not be too difficult to find at least one smell that is easily recognizable and identifiable within each of these six smell categories, which incidentally also has a fair distribution ratio between what can be termed as pleasant and unpleasant smell categories. Moreover, as this smell categorization scheme has a small number of categories compared to other classification schemes [5], [9], odor adaptation issues, such as losing one's sensitivity to detect smells after continuous exposure and prolonged stimulation of the olfactory sense, are reduced to a minimum. The six smell categories used were *flowery, foul, fruity, burnt, resinous* and *spicy*. Table I shows a brief summary of the videos

TABLE II
IMPACT OF DELAY ON THE USER-PERCEIVED EXPERIENCE

Questionnaire Statement	-30	-20	-10	0	+10	+20	+30
• The smell was relevant to what I was watching	✕	✕	✕	✕	✕	✕	✕
• The smell heightened the sense of reality whilst watching the video clip	✕	✕	✕	✕	✕	✕	✓
• The smell was distracting	✕	✕	✕	✕	✕	✕	✕
• The smell was annoying	✕	✕	✕	✕	✕	✕	✕
• I enjoyed watching the video clip	✕	✕	✕	✕	✕	✕	✕

used for each of the smell categories and what smell was generated during its playback.

B. Experimental Participants

A total of 42 participants, 14 females and 28 males, between the ages of 18 and 40, from a wide variety of backgrounds, i.e., undergraduate and postgraduate students from different universities and departments, and both blue and white collar workers, took part in the experimental study.

C. Scented Audiovisual Synchronization Skews

We define scented audiovisual synchronization as the temporal relationship between a computer generated stream of smell and audiovisual streams. Perfect synchronization between the olfactory and audiovisual media streams means that the desired temporal relationship is achieved during the olfaction-enhanced multimedia presentation. In other words, the olfactory and audiovisual media streams have zero skew between them, as they are in sync. Our study was performed by introducing artificially created skews at intervals of 10 s ranging from −30 s to +30 s between olfactory stimuli and audiovisual media content. This was achieved by programming the Vortex Active to start the release of the smell at the skewed start time in relation to its synchronization with the video playback.

D. Experimental Process

Each experimental case lasted for approximately 20 min with participants watching the set of six olfactory enhanced video clips of 240×180 pixels as described in Table I. Each video clip lasted for about 90 s, with the duration of the smell lasting for about 30 s. The excerpts of the video clips were specifically chosen so that audiovisual content relating to the olfactory data the videos were augmented with coincided with the middle 30 s segment. Thus in this respect, a 90 s olfaction-enhanced multimedia clip may be thought of comprising the following 3 synchronization temporal segments: the first 30 s representing olfaction ahead of audiovisual content, the middle 30 s represent olfaction "in sync," and the last 30 s representing olfaction behind audiovisual content.

V. RESULTS

We present an analysis of participants' responses to the questions asked at the end of each video clip watched. A significance level of $p < 0.05$ was adopted for the study.

A. Acceptable Inter-Media Skew Between Olfactory and Audiovisual Media Content

To determine scented audiovisual synchronization requirements, i.e., the acceptable boundaries within which the presentation of the olfactory data in relation to audiovisual content is perceived as being timely by users, participants were asked to respond to the statement: *The smell*
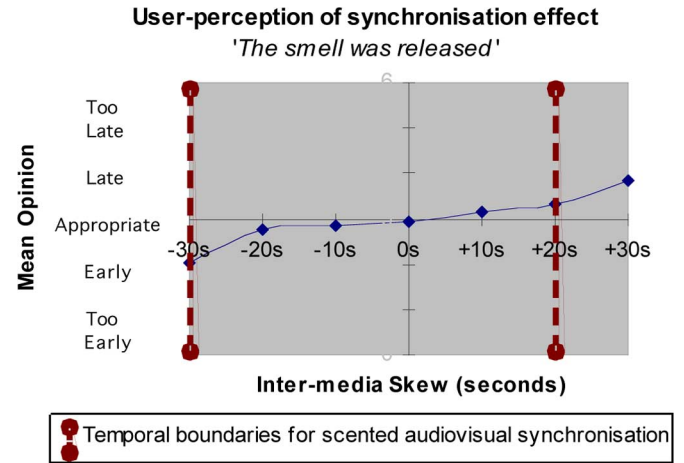

Fig. 2. Overall analysis of synchronization errors with respect to applied inter-media skew between olfactory data and audiovisual content.

was released. Response options on a rating scale of 1 to 5 were: "*Too Early*," "*Early*," "*At an Appropriate Time*," "*Late*," and "*Too Late*."

Our results showed that significant differences in participants' perception of the scented audiovisual synchronization effect occurred at an inter-media skew of −30 s ($\text{Mean} = 2.04, \text{t} = -4.830, \text{p} < 0.05$), i.e., olfaction ahead of audiovisual content by 30 s or more, and at an inter-media skew value greater than +20 s ($\text{Mean} = 3.32, \text{t} = -2.436, \text{p} < 0.05$ at +20 s and $\text{Mean} = 3.83, \text{t} = -4.889, \text{p} < 0.05$ at +30 s), i.e., when olfaction is behind audiovisual content by more than 20 s. At each of the other artificially introduced inter-media skew values ranging between −20 s to +10 s, the differences in participants' perception of the synchronization effect were not found to be statistically significant when compared with the case for perfect synchronization. Bearing in mind that skews in our experiment ranged from −30 s to 30 s, our results reveal two temporal boundary regions for synchronizing olfactory media with video media (Fig. 2).

- **The In-Sync Region**: where this region spans between a maximum skew of −30 s, when olfaction is ahead of audiovisual, and a maximum skew of +20 s for when olfaction is behind audiovisual. That is, olfactory media content may be generated up to 30 s ahead of audiovisual content or delayed by a maximum of 20 s before users begin to notice that there is a synchronization error present. Accordingly, these temporal boundaries represent the synchronization requirements when olfactory media content is combined with audiovisual content.
- **An Out-Of-Sync Region**: this region spans beyond +20 s, when olfaction is behind audiovisual content. Here, the majority of participants detected the synchronization error.

Furthermore, the temporal boundaries established show that olfaction behind audiovisual content is more noticeable than the case when olfaction is ahead. This is reflected by the slightly smaller range at

Fig. 3. Olfaction-enhanced multimedia—Impact.

Fig. 4. Perceived olfactory synchronization.

which the presence of inter-media skew becomes noticeable when olfaction is behind, i.e., an inter-media skew range of 20 s, compared to an inter-media skew range of 30 s when olfaction is ahead. Nonetheless, this is consistent with the nature of smell to sometimes linger and drift when emitted, with the consequence that such a natural delay occurring in the presence of inter-media skew involving olfactory media content will definitely favor the case for when olfaction is ahead of other media content.

B. Impact of Delay on the User-Perceived Experience

In this section, we examine the impact that the synchronization error had on participants' experience of olfaction-enhanced multimedia video excerpts. The influence of synchronization errors on the user-perceived experience will be ascertained according to participants' opinions in respect of the following:

- user-perceived experience of the sense of relevance of the olfactory media;
- user-perceived experience of the sense of reality of the olfaction-enhanced multimedia;
- user acceptability of olfactory media in terms of how distracting or annoying they find it;
- user-perceived enjoyment of the olfaction-enhanced multimedia experience.

In Table II and Fig. 3, we present a summary of an independent samples t-test analysis used to compare participants' opinions of the user-perceived experience of the olfactory-enhanced videos in the presence of skew between the olfactory and audiovisual media content with the case of perfect synchronization. Our results revealed that generally there was no significant difference in the user-perceived experience of the sense of relevance and reality of the olfactory media, the level of annoyance or distraction it created or the overall sense of enjoyment of the olfaction-enhanced multimedia experience in the presence and absence of skew. The only exception to this was in the case of the user-perceived experience of the sense of reality at a skew of $+30$ s ($\text{Mean} = 3.23, \text{t} = -2.237, \text{p} < 0.05$), i.e., olfaction 30 s behind audiovisual content.

This exception for the user-perceived experience of the sense of reality in the presence of a skew of $+30$ s is, however, not surprising, considering that a skew of $+30$ s occurs at the tail end of the last segment of the multimedia clip. Each clip comprised of three segments of 30 s, representing olfaction ahead, olfaction in-sync and olfaction behind, respectively. Unlike the perceived sense of relevance, for example, in which case participants were probably able to cast their minds back to what was watched to relate the olfactory media to the video content, the

same approach would have been futile with respect to the perception of reality. The sense of reality is more concerned with present timing, and not in the past or future.

Last, but not least, in as far as the question of whether users judge different smells differently, an analysis of variance tests highlights (Fig. 4) that, for each olfactory skew considered in our work, there are no significant differences in the user enjoyment of the olfactory-enhanced multimedia presentation.

VI. CONCLUDING DISCUSSION

Results reported in this paper showed that significant differences of participants' perception of the synchronization effect between olfactory media and video content occurred at inter-media skew values of -30 s, for olfaction ahead of audiovisual, and $+20$ s and $+30$ s, for when olfaction was behind audiovisual content. Consequently, we were able to conclude that the temporal boundaries for creating a scented audiovisual experience that will be perceived as being synchronised by users lie within a temporal range of -30 s to $+20$ s. Accordingly, we propose that the inter-media skew requirements for combining video with olfaction are -30 s (olfaction ahead of video) and $+20$ s (video ahead of olfaction). These inter-media skew synchronization requirements, as proposed here, are applicable to the creation of scented audiovisual experiences for users, e.g., scented videos and multi-sensory virtual reality experiences.

Furthermore, our results also revealed that although participants detected the presence of synchronization errors, it did not have a significant impact on the general perceptual experience of the olfaction-enhanced multimedia for participants. Participants' opinions still reflected positive attitudes with regards to the enjoyment of the multimedia experience, the sense of relevance of the olfactory media and the sense of reality they felt it created, both in the absence and presence of inter-media skew. The slight exception to this general opinion concerns participants' perception of the sense of reality of the olfaction-enhanced multimedia. In this case, participants' opinions reflected that at a skew of $+30$ s (olfaction behind video), the olfactory media no longer had the impact to heighten the sense of reality of the multimedia experience. In addition, participants generally did not find the delayed or advance emission of scent distracting or annoying, and their opinions mostly reflected a negative bias or neutral opinion in this respect, both in the absence or presence of inter-media skew.

We recognize, though, two limitations of our study: our investigation only looked at skew values within the $[-30$ s$; +30$ s$]$ interval; secondly, the set of parameters that one can hold constant when conducting

user perceptual experiments is much smaller in the case of olfactory data than when dealing with traditional media such as video and audio. While the latter will always be a challenge for olfactory-enhanced multimedia research, the former presents opportunity for valuable future work.

REFERENCES

[1] G. Blakowski and R. Steinmetz, "A media synchronization survey: Reference model, specification, and case studies," *IEEE J. Select. Areas Commun.*, vol. 14, no. 1, pp. 5–35, Jan. 1996.

[2] A. Bodnar, R. Corbett, and D. Nekrasovski, "AROMA: Ambient awareness through olfaction in a messaging application," in *Proc. 6th Int. Conf. Multimodal Interfaces*, State College, PA, 2004, pp. 183–190.

[3] S. A. Brewster, D. K. McGookin, and C. A. Miller, "Olfoto: Designing a smell-based interaction," in *Proc. Conf. Human Factors in Computing Systems*, Montréal, QC, Canada, 2006, pp. 653–662.

[4] V. Cevher, A. C. Sankaranarayanan, J. H. McClellan, and R. Chellappa, "Target tracking using a joint acoustic video system," *IEEE Trans. Multimedia*, vol. 9, no. 4, pp. 715–727, Jun. 2007.

[5] M. Chastrette, "Classification of odors and structure-odor relationships," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 100–116.

[6] R. S. Herz, "Influences of odors on mood and affective cognition," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 160–177.

[7] R. Hudson and H. Distel, "The individuality of odor perception," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 408–420.

[8] T. Hummel, S. Heilmann, and C. Murphy, "Age-related Changes in Chemosensory Functions," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 441–456.

[9] N. Kaye, "Symbolic olfactory display," M.S. thesis, Massachusetts Inst. Technol., Cambridge, MA, 2001. [Online]. Available: http://www.media.mit.edu/~jofish/thesis/.

[10] A. Keller and L. B. Vosshall, "Human olfactory psychophysics," *Current Biol.*, vol. 14, no. 20, pp. R875–R878, 2004.

[11] E. P. Köster, "The Specific Characteristics of the Sense of Smell," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 27–44.

[12] S. Saito, S. Ayabe-Kanamura, Y. Takashima, N. Gotow, N. Naito, T. Nozawa, M. Mise, Y. Deguchi, and T. Kobayakawa, "Development of a smell identification test using a novel stick-type odor presentation kit," *Chem. Senses*, vol. 31, no. 4, pp. 379–391, 2006.

[13] M. E. Sargin, Y. Yemez, E. Erzin, and A. M. Tekalp, "Audiovisual synchronization and fusion using canonical correlation analysis," *IEEE Trans. Multimedia*, vol. 9, no. 7, pp. 1396–1403, Nov. 2007.

[14] B. Schaal, R. Soussignan, and L. Marlier, "Olfactory Cognition at the Start of Life: The Perinatal Shaping of Selective Odor Responsiveness," in *Olfaction, Taste, and Cognition*, C. Rouby, B. Schaal, D. Dubois, R. Gervais, and A. Holley, Eds., Paperback Re-issue ed. New York: Cambridge Univ. Press, 2002, pp. 421–440.

[15] R. Steinmetz, "Human perception of jitter and media synchronization," *IEEE J. Select. Areas Commun.*, vol. 14, no. 1, pp. 61–72, Jan. 1996.

Feature Selection Under a Complexity Constraint

Jan H. Plasberg and W. Bastiaan Kleijn, *Fellow, IEEE*

Abstract—**Classification on mobile devices is often done in an uninterrupted fashion. This requires algorithms with gentle demands on the computational complexity. The performance of a classifier depends heavily on the set of features used as input variables. Existing feature selection strategies for classification aim at finding a "best" set of features that performs well in terms of classification accuracy, but are not designed to handle constraints on the computational complexity. We demonstrate that an extension of the performance measures used in state-of-the-art feature selection algorithms with a penalty on the feature extraction complexity leads to superior feature sets if the allowed computational complexity is limited. Our solution is independent of a particular classification algorithm.**

Index Terms—**Classification, complexity, context awareness, cost, feature selection, mutual information.**

I. INTRODUCTION

The emergence of ubiquitous computing is fueled by an increase in the number of affordable portable devices. Future devices will be more intelligent and are expected to be context-aware, e.g., a mobile phone senses a meeting and switches to quiet mode. However, for continuous monitoring of the context on mobile devices it is important to constrain the computational effort [1]. Moreover, the design engineer typically has the task to develop solutions using limited resources.

Classification accuracy is strongly dependent on the set of features used as input variables. For automatic feature selection, the literature provides us with numerous strategies aiming to find a "best" set of features (for an overview, see, for example, [2]–[6]). In this paper we aim at finding a "best" set of features given a constraint on the computational complexity or cost of the feature acquisition, which may dominate the cost of the classifier. The cost of feature acquisition has been considered in [7]–[14]. The solution in [7] does not readily translate into large problem sizes involving hundreds of features. All other cited methods suffer from one or more of the following restrictions: the cost and the benefit of adding a feature have to be in the same measure, ad hoc trade-offs between cost and benefit are made or the feature selection is dependent on the classification algorithm used.

This paper is an extension to our previous work [15]. In this paper, we present a Lagrangian approach to extend the optimization criteria used by sequential feature selection strategies with an additional cost incorporating the computational complexity induced by choosing a particular feature. We allow the cost and the benefit of a feature to be in different measures. Our solution is independent of a particular classification algorithm and leads to feature sets that fulfill strict constraints on the total cost of feature acquisition. While the cost of feature acquisition is the main topic of the paper, we are also concerned with the computational requirements posed by the feature selection algorithm and clarify whenever these two measures might get confused.