

fastGCVM: A Fast Algorithm for the Computation of the Discrete Generalized Cramér-von Mises Distance

Johannes Meyer
Karlsruhe Institute of
Technology,
Vision and Fusion
Laboratory,
Adenauerring 4
76131 Karlsruhe,
Germany
johannes.meyer@kit.edu,
www.meyer-research.de

Thomas Längle
Fraunhofer Institute of
Optronics,
System Technologies and
Image Exploitation IOSB,
Fraunhoferstr. 1
76131 Karlsruhe,
Germany
thomas.laengle
@iosb.fraunhofer.de

Jürgen Beyerer
Fraunhofer Institute of
Optronics,
System Technologies and
Image Exploitation IOSB,
Fraunhoferstr. 1
76131 Karlsruhe,
Germany
juergen.beyerer
@iosb.fraunhofer.de

ABSTRACT

Comparing two random vectors by calculating a distance measure between the underlying probability density functions is a key ingredient in many applications, especially in the domain of image processing. For this purpose, the recently introduced generalized Cramér-von Mises distance is an interesting choice, since it is well defined even for the multivariate and discrete case. Unfortunately, the naive way of computing this distance, e.g., for two discrete two-dimensional random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in [0, \dots, n-1]^2, n \in \mathbb{N}$ has a computational complexity of $O(n^5)$ that is impractical for most applications. This paper introduces fastGCVM, an algorithm that makes use of the well known concept of summed area tables and that allows to compute the generalized Cramér-von Mises distance with a computational complexity of $O(n^3)$ for the mentioned case. Two experiments demonstrate the achievable speed up and give an example for a practical application employing fastGCVM.

Keywords

Distance of random vectors, summed area tables, speed up, histogram comparison, localized cumulative distributions, generalized Cramér-von Mises distance.

1 INTRODUCTION

Applications and algorithms from different fields require to compute a distance between two random vectors, respectively, between the corresponding probability density functions in order to measure their similarity [Cha07]. For example, histogram distances are employed by content based image retrieval systems to find images similar to the query image [MGW10, CS02, DNK03]. Furthermore, such distance measures can be used as an optimization criterion to obtain a Dirac mixture approximation of probability distributions, for interpolations, for parameter estimation or for tracking [PHB13]. Whenever one of the two random vectors is of discrete type, the cumulative distribution functions are usually employed for further processing steps. However, this is only possible for the one-dimensional

case since the cumulative distribution of a multivariate random vector is not unique.

In [HK08], the authors have introduced a novel formulation of the cumulative distribution, the so-called localized cumulative distribution. The LCD of a random vector can be imagined as a rectangular kernel transform of the underlying probability density function. Based on the definition of the LCD, [HK08] introduces a generalized formulation of the Cramér-von Mises (CVM) distance that can be used to calculate distances between two (multivariate) random vectors of whom one (or both) may even be of discrete type. However, as it is shown in the following sections, the computation of the LCD and the CVM distance for two discrete two-dimensional random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in [0, \dots, n-1]^2, n \in \mathbb{N}$, as it is typical for image processing applications, has a complexity of $O(n^5)$.

This paper shows, how the concept of so-called summed area tables [Cro84, VJ04] can be used to obtain a novel and more efficient algorithm for computing the LCD and the CVM distance. For the mentioned example random vectors, the proposed algorithm has a reduced complexity of $O(n^3)$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

The paper is structured as follows: Sec. 2 lists and discusses related work performed by other research groups. In Sec. 3, the localized cumulative distributions and the generalized Cramér-von Mises distance are introduced and their computational complexity is analyzed. Furthermore, the concept of summed area tables is described. Section 4 is dedicated to the proposed fast algorithm for the computation of the Cramér-von Mises distance and Sec. 5 covers the performed experiments and the respective results. A summary of the paper and an outlook concerning further research topics is provided in Sec. 6.

2 RELATED WORK

This section describes major work performed by other researches in which either the generalized Cramér-von Mises distance plays an important role or summed area tables have been used. To the knowledge of this paper's authors, neither the concept of summed area tables nor any other acceleration technique has yet been employed to reduce the computation time of the generalized Cramér-von Mises distance.

Franklin Crow was the first to introduce the concept of summed area tables [Cro84]. The respective paper is concerned with the task of reducing the computational costs of the texture mapping problem. Crow shows that it is possible to precompute a certain data structure, a summed area table (see Sec. 3.1), for a given image using linear time and space so that afterwards the sum of the pixel values inside any arbitrary query rectangle can be obtained in constant time.

In [VJ04], Viola and Jones introduce a processing framework for face detection. They employ a large set of features in concert with a classifier cascade trained using the AdaBoost learning algorithm. By adapting the idea of summed area tables to digital images, they can achieve a fast computation of the used features.

Hanebeck et al. and Gilitschenski et al. show in [HHK09, GH13], how the concept of localized cumulative distributions and the generalized Cramér-von Mises distance can be employed to obtain Dirac mixture approximations of multivariate Gaussian distributions. Such approximations are important ingredients, e.g., for state estimation in dynamic systems. In their presented work, the authors achieve an efficient implementation by analytically obtaining a closed-form solution for the LCD and the generalized CVM distance for multivariate Gaussian densities and Dirac mixtures. Since the approach for accelerating the calculation of the LCD and the CVM distance provided in this paper does not rely on any specific form of a probability distribution, it can be used to speed up applications, where no analytical solution can be found.

3 PREREQUISITES

Before the details of the acceleration approach can be described, this section introduces the necessary definitions and data structures beginning with the concept of summed area tables. For the sake of simplicity, the formulas and solutions shown in this paper are often limited to the case of discrete two-dimensional probability distributions since this is the common case for image processing applications.

3.1 Summed Area Tables

Summed area tables denote a data structure that can be precomputed for two-dimensional input arrays allowing to calculate the sum of the array entries inside arbitrary rectangular regions of the array [Cro84].

Let $i(x, y) \in \mathbb{R}$, $x, y \in [0, \dots, n-1]$ denote a two-dimensional array like data structure. The corresponding summed area table I is defined by

$$I(x, y) := \begin{cases} 0 & \text{if } \min\{x, y\} \leq 0, \\ \sum_{x_f=1}^x \sum_{y_f=1}^y i(x_f, y_f) & \text{otherwise.} \end{cases} \quad (1)$$

The data structure I can be calculated in a single sweep over i by employing the iterative formulation

$$I(x, y) = i(x, y) + I(x-1, y) + I(x, y-1) - I(x-1, y-1). \quad (2)$$

By this means, the sum of array entries of i inside an interval $x \in [x_f, x_t]$, $y \in [y_f, y_t]$ can be obtained via three arithmetic operations on I only:

$$\begin{aligned} \sum_{x=x_f}^{x_t} \sum_{y=y_f}^{y_t} i(x, y) = & I(x_t, y_t) \\ & - I(x_f - 1, y_t) \\ & - I(x_t, y_f - 1) \\ & + I(x_f - 1, y_f - 1). \end{aligned} \quad (3)$$

Figure 1 provides a visualization motivating the formula. Since the iterative expression (2) involves only four constant time arithmetic operations and four constant time array accesses and has to be performed for every element in i , i.e., $n \cdot n = n^2$ times, the computation of I has a complexity of $O(n^2)$ [MS08, Cor09]. By using formula (3), the calculation of a sum of array entries of i in an arbitrary rectangular region has a complexity of $O(1)$. The concept of summed area tables is used in later sections in order to accelerate the computation of the generalized Cramér-von Mises distance.




Figure 1: Calculating the sum of the array entries inside the query rectangle $x \in [x_f, x_t]$, $y \in [y_f, y_t]$ (purple) by employing a summed area table: from the red component $I(x_t, y_t)$, first the orange component $I(x_f - 1, y_t)$ and the blue component $I(x_t, y_f - 1)$ have to be subtracted. Since the area represented by the green component has now been subtracted twice, $I(x_f - 1, y_f - 1)$ has to be added.

3.2 Localized Cumulative Distribution

As mentioned before, cumulative distributions are often employed in order to be able to calculate a distance between two random vectors. However, the conventional definition of the cumulative distribution is not unique for the multivariate case. This issue is tackled by the so-called localized cumulative distribution LCD introduced in [HK08]. For a given random vector $\tilde{\mathbf{x}} \in \mathbb{R}^N$, $N \in \mathbb{N}$ and the corresponding probability density function $f: \mathbb{R}^N \rightarrow \mathbb{R}_+$, the respective LCD $F(\mathbf{x}, \mathbf{b})$ is defined as

$$F(\mathbf{x}, \mathbf{b}) := P\left(|\tilde{\mathbf{x}} - \mathbf{x}| \leq \frac{1}{2}\mathbf{b}\right), \quad (4)$$

$$F(\cdot, \cdot): \Omega \rightarrow [0, 1], \Omega \subset \mathbb{R}_+^N \times \mathbb{R}_+^N,$$

$$\mathbf{b} \in \mathbb{R}_+^N,$$

with $\mathbf{x} \leq \mathbf{y}$, $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^N$ denoting a component-wise relation that only holds if $\forall j \in [1, \dots, N]: x_j \leq y_j$. The LCD $F(\mathbf{x}, \mathbf{b})$ can be imagined as an integral transform with the rectangular kernel \mathbf{b} providing the limits of the integration for the different dimensions. Based on the

probability density function $f(\mathbf{x})$ corresponding to $\tilde{\mathbf{x}}$, the respective LCD $F(\mathbf{x}, \mathbf{b})$ is calculated by

$$F(\mathbf{x}, \mathbf{b}) = \begin{cases} \int_{\mathbf{x} - \frac{1}{2}\mathbf{b}}^{\mathbf{x} + \frac{1}{2}\mathbf{b}} f(\mathbf{t}) d\mathbf{t} & , \text{ if } \tilde{\mathbf{x}} \text{ continuous,} \\ \sum_{\mathbf{t}=\max\{\mathbf{0}, \mathbf{x} - \lfloor \frac{1}{2}\mathbf{b} \rfloor\}}^{\min\{\mathbf{x}_{\max}, \mathbf{x} + \lfloor \frac{1}{2}\mathbf{b} \rfloor\}} f(\mathbf{t}) & , \text{ if } \tilde{\mathbf{x}} \text{ discrete,} \end{cases} \quad (5)$$

with $\mathbf{0} = (0, \dots, 0)^T$ denoting the zero vector, \mathbf{x}_{\max} denoting the vector representing the upper limit of the spatial support and $\max\{\mathbf{x}\}$, $\min\{\mathbf{x}\}$ and $\lfloor \mathbf{x} \rfloor$ denoting elementwise operations. In contrast to the conventional definition of the cumulative distribution function, the LCD of a multivariate random vector is unique. Since this paper is especially focused on the case of discrete random vectors, the discrete case of Eq. (5) will be considered in further sections.

3.2.1 Complexity of localized cumulative distribution evaluation

In order to determine the computational complexity of one evaluation of the LCD $F(\mathbf{x}, \mathbf{b})$ for a given probability density function of a discrete random vector $\tilde{\mathbf{x}} \in \mathbb{R}^N$, the kernel vector \mathbf{b} is considered to hold the same value in all dimensions, i.e., $\mathbf{b} = (b, \dots, b)^T$, what is always the case for its usage in the context of the generalized Cramér-von Mises distance. Since evaluating $F(\mathbf{x}, \mathbf{b})$ requires b array accesses and summation operations along each dimension, the corresponding complexity is $O(b^N)$ and hence $O(b^2)$ for the two-dimensional case common in image processing applications.

3.3 Generalized Cramér-von Mises Distance

Based on the introduced localized cumulative distribution (4), the generalized Cramér-von Mises distance between two multivariate random vectors can now be defined employing their LCDs as shown in [HK08]. For the LCDs $F(\mathbf{x}, \mathbf{b})$, $G(\mathbf{x}, \mathbf{b})$ corresponding to the probability density functions $f(\mathbf{x})$, $g(\mathbf{x})$ of two random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in \mathbb{R}^N$, $N \in \mathbb{N}$, their generalized Cramér-von Mises distance is given by

$$D(f, g) := \int_{\mathbb{R}_+^N} \int_{\mathbb{R}_+^N} (F(\mathbf{x}, \mathbf{b}) - G(\mathbf{x}, \mathbf{b}))^2 d\mathbf{b} d\mathbf{x}. \quad (6)$$

In the case of discrete random vectors, the integrals are replaced by summations resulting in:

$$D(f, g) = \sum_{\mathbf{x} \in \Omega_s} \sum_{b=0}^{b_{\max}} (F(\mathbf{x}, (b, \dots, b)^T) - G(\mathbf{x}, (b, \dots, b)^T))^2, \quad (7)$$

with Ω_s denoting the spatial support of the probability density functions and b_{\max} representing the absolute maximum component value of Ω_s , i.e., the maximum kernel size necessary to capture the whole probability density function.

3.3.1 Complexity of generalized Cramér-von Mises distance calculation

This section deals with the computational complexity of the calculation of the discrete CVM distance (7). As it is the most important case for image processing applications, the CVM distance for two discrete two-dimensional random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in [0, \dots, n-1]^2$ and the corresponding probability density functions f, g is considered and the resulting complexity is determined in successive steps. As shown in Sec. 3.2.1, one evaluation of the LCD $F(\mathbf{x}, \mathbf{b})$ has a complexity of $O(b^2)$ and thus every loop of the inner summation of Eq. (7) also has a complexity of $O(b^2)$. Since b is increased by 1 from 0 to $n-1$, the computation of the whole inner sum of Eq. (7) requires a number of

$$\sum_{b=0}^{n-1} b^2 = \frac{n^3}{3} - \frac{n^2}{2} + \frac{n}{6} \quad (8)$$

computations and hence has a complexity of $O(n^3)$. Conclusively, as all these computations have to be performed over the whole spatial support of the underlying probability density functions, i.e., a total of n^2 times, the overall complexity of the computation of $D(f, g)$ is in $O(n^5)$:

$$D(f, g) = \underbrace{\sum_{\mathbf{x} \in \Omega_s} \underbrace{\sum_{b=0}^{n-1} \overbrace{(F(\mathbf{x}, (b, b)^T) - G(\mathbf{x}, (b, b)^T))^2}_{\in O(b^2)}}_{\in O(n^3)}}_{\in O(n^5)} \quad (9)$$

4 FAST GENERALIZED CRAMÉR-VON MISES DISTANCE

For the case of two two-dimensional discrete random vectors, the calculation of the generalized CVM distance can be accelerated. Therefore, the concept of summed area tables is employed in the evaluation of the localized cumulative distribution (5). Since for the referenced case, equation (5) represents a summation over a rectangular region, a summed area table can be precomputed for $f(\mathbf{x})$ in order to speed up the evaluation of $F(\mathbf{x}, \mathbf{b})$. Algorithm 1 lists the pseudo code for the proposed fast calculation of the generalized CVM distance. After building the summed area tables $\mathfrak{f}, \mathfrak{g}$ for the discrete input probability density functions f and

g , the squared difference between the localized cumulative distributions of f and g are computed for every spatial position $(i, j)^T$ and for every sensible kernel size $b \in [0, \dots, n-1]$. The LCDs are evaluated by Algorithm 2 that employs Eq. (3) in concert with the provided summed area table \mathfrak{s} , the spatial position $(i, j)^T$ and the kernel size $(b, b)^T$ to obtain the summation of the values of the probability density function inside the respective rectangle.

Algorithm 1 Fast algorithm for calculating the generalized Cramér-von Mises distance for two two-dimensional discrete random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}}$ represented by their corresponding probability density functions $f(\mathbf{t}), g(\mathbf{t})$.

```

fastGCVM( $f(\mathbf{x}), g(\mathbf{x})$ )
 $\mathfrak{f} \leftarrow$  generateSummedAreaTable( $f(\mathbf{x})$ )
 $\mathfrak{g} \leftarrow$  generateSummedAreaTable( $g(\mathbf{x})$ )
 $D \leftarrow 0$ 
for  $i = 0, \dots, n-1$  do
  for  $j = 0, \dots, n-1$  do
    for  $b = 0, \dots, n-1$  do
       $D \leftarrow D +$ 
        ( $\text{fastLCD}(\mathfrak{f}, (i, j)^T, (b, b)^T) -$ 
 $\text{fastLCD}(\mathfrak{g}, (i, j)^T, (b, b)^T)$ )2
    end for
  end for
end for
return  $D$ 

```

Algorithm 2 Algorithm for evaluating the localized cumulative distribution $LCD((i, j)^T, (b, b)^T)$ of a two-dimensional discrete random vector at position $(i, j)^T$ with kernel sizes $(b, b)^T$ based on the summed area table \mathfrak{s} of the underlying probability density function.

```

fastLCD( $\mathfrak{s}, (i, j)^T, (b, b)^T$ )
 $x_f \leftarrow \max\{0, i - \lfloor \frac{1}{2}b \rfloor\}$ 
 $x_t \leftarrow \min\{x_{\max}, i + \lfloor \frac{1}{2}b \rfloor\}$ 
 $y_f \leftarrow \max\{0, j - \lfloor \frac{1}{2}b \rfloor\}$ 
 $y_t \leftarrow \min\{y_{\max}, j + \lfloor \frac{1}{2}b \rfloor\}$ 
return  $\mathfrak{s}(x_t, y_t) - \mathfrak{s}(x_f - 1, y_t) - \mathfrak{s}(x_t, y_f - 1)$ 
 $+ \mathfrak{s}(x_f - 1, y_f - 1)$ 

```

4.1 Complexity of fastGCVM Algorithm

The fastGCVM algorithm shown in the listing Algorithm 1 has three nested loops with n iterations each. The inner loop calculates a squared difference between the results of two evaluations of the fastLCD algorithm shown in listing Algorithm 2. Since fastLCD requires only a constant number of max, min, array lookups and arithmetic operations, it has a constant complexity of $O(1)$. Consequently, the complexity of the inner loop

Table 1: Execution times resulting from the performed experiment. The size of the input histograms is denoted by n and t_{naiv} , t_{fast} represent the mean measured execution times in milliseconds \pm standard deviation for the naive, respectively, the fast algorithm.

| $n =$ | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 |
|----------------------------|--------------------|--------------------|--------------------|--------------------|--------------------|---------------------|----------------------|------------------------|------------------------|----------------------|
| t_{naiv} in ms | 7.2 ± 0.6 | 183.5 ± 0.3 | 1,311 ± 2 | 5,379 ± 4 | 16,145 ± 36 | 49,379 ± 138 | 105,820 ± 218 | 201,940 $\pm 1,550$ | 375,010 $\pm 1,433$ | 634,180 ± 390 |
| t_{fast} in ms | 0.23 ± 0.02 | 0.75 ± 0.01 | 2.18 ± 0.01 | 4.96 ± 0.02 | 9.7 ± 0.3 | 16.29 ± 0.06 | 25.7 ± 0.1 | 38.35 ± 0.08 | 54.6 ± 0.2 | 75.1 ± 0.1 |

of Algorithm 1 is in $O(1)$ too. As the three loops run n iterations each, the algorithm fastGCVM has a total complexity of $O(n^3)$.

5 EXPERIMENTS

Two experiments have been designed in order to demonstrate the effectiveness of the proposed speed-up technique. In the first experiment, multiple pairs of two-dimensional histograms with numbers of bins ranging from 10^2 to 100^2 have been randomly generated and used as the discrete input probability density functions. Between each pair, the generalized Cramér-von Mises distance has been calculated using a naive implementation of Eq. 7 and the proposed fastGCVM method shown in listing Algorithm 1. The algorithms have been implemented in C# using the Accord.NET framework [Sou14] and by not making use of any parallelism. The measured execution times are listed in Table 1. The results clearly show that the fast algorithm fastGCVM outperforms the naive implementation even for the case of small problem instances. For the largest employed example histogram with $100 \times 100 = 10,000$ bins, the naive implementation requires more than 10 minutes for calculating the desired result whereas fastGCVM needs only about 75 milliseconds.

For the second experiment, the code of an existing application has been extended by the fastGCVM algorithm. In [MLB16a, MLB16b], so-called light deflection maps are processed in order to visually inspect transparent objects for material defects. Deflection maps are spatially resolved data structures similar to discrete histograms containing information about the angles by which light rays get deflected while propagating through a transparent test object. Strong spatial discontinuities between adjacent deflection maps provide an indication of present scattering material defects that lead to changes in the distribution of the light's propagation direction. Such discontinuities cause high values of the generalized Cramér-von Mises distance between spatially adjacent deflection maps. In the context of the second experiment, the execution time was measured that has been required for processing the deflection maps [MLB16a] of a transparent cylindrical lens using both the naive and the fast implementation. For each input data set, the generalized CVM distance had

to be calculated 3042 times between histograms having $9 \times 9 = 81$ bins. Figure 2 shows the resulting inspection images in pseudo colors. The naive implementation of the generalized CVM had an execution time of $15,750 \text{ ms} \pm 2 \text{ ms}$ and the fastGCVM algorithm finished after $180 \text{ ms} \pm 2 \text{ ms}$.

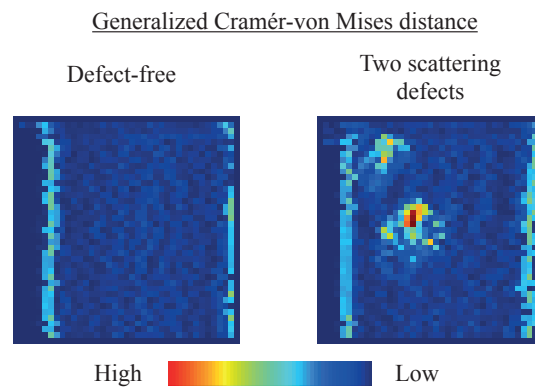


Figure 2: Pseudo color inspection images for plano-convex cylindrical lenses resulting from calculating the generalized Cramér-von Mises distance for spatially adjacent deflection maps. The left image corresponds to a defect-free test object instance and the right image corresponds to a test object instance affected by two scattering surface defects. The two defects are clearly indicated by the two regions of higher intensities in the image's upper left corner.

In summary, the experiments show that using the naive implementation of the Cramér-von Mises distance is not suitable for any practical application where execution time plays a critical role. Especially for visual inspection systems—as shown in the second example—which often have to fulfill real-time requirements, the fastGCVM algorithm allows to employ the generalized Cramér-von Mises distance in the image processing pipeline due to its reduced computational complexity.

6 SUMMARY

The generalized Cramér-von Mises distance is a helpful tool for comparing multivariate random vectors. However, for the frequent case of two discrete two-dimensional random vectors $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in [0, \dots, n-1]^2$, $n \in \mathbb{N}$, the naive implementation of the CVM distance

has a computational complexity of $O(n^5)$, what leads to execution times impractical for many applications. After introducing the reader to the idea of summed area tables, which is a common speed up technique in the domain of image processing, the generalized Cramér-von Mises distance and the underlying concept of localized cumulative distributions have been introduced. The paper then proposes to employ summed area tables in order to obtain fastGCVM, a fast algorithm for calculating the generalized Cramér-von Mises distance with a reduced computational complexity of only $O(n^3)$. By means of two experiments it could be shown that fastGCVM clearly outperforms the naive implementation of the generalized CVM distance—in some cases, fastGCVM's execution time is four magnitudes lower than the time required by the naive implementation. It should be mentioned, that the execution time of fastGCVM can be further reduced by adequately employing simple parallelization techniques. As further steps, the authors plan to provide their implementation as an open source library and to theoretically show, how summed area tables can be used to also speed up the calculation of the generalized CVM for higher dimensional discrete random vectors.

7 REFERENCES

- [Cha07] Sung-Hyuk Cha. Comprehensive survey on distance/similarity measures between probability density functions. *City*, 1(2):1, 2007.
- [Cor09] Thomas H Cormen. *Introduction to algorithms*. MIT press, 2009.
- [Cro84] Franklin C. Crow. Summed-area tables for texture mapping. *ACM SIGGRAPH computer graphics*, 18(3):207–212, 1984.
- [CS02] Sung-Hyuk Cha and Sargur N. Srihari. On measuring the distance between histograms. *Pattern Recognition*, 35(6):1355–1370, 2002.
- [DNK03] Dietrich Van der Weken, Mike Nachtegael, and Etienne Kerre. Using similarity measures for histogram comparison. In *Fuzzy Sets and Systems-IFSA 2003*, pages 396–403. Springer, 2003.
- [GH13] Igor Gilitschenski and Uwe D. Hanebeck. Efficient deterministic dirac mixture approximation of Gaussian distributions. pages 2422–2427. IEEE, 2013.
- [HHK09] Uwe D. Hanebeck, Marco F. Huber, and Vesa Klumpp. Dirac mixture approximation of multivariate gaussian densities. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 3851–3858. IEEE, 2009.
- [HK08] Uwe D. Hanebeck and Vesa Klumpp. Localized Cumulative Distributions and a Multivariate Generalization of the Cramér-von Mises Distance. In *Multisensor Fusion and Integration for Intelligent Systems, 2008. MFI 2008. IEEE International Conference on*, pages 33–39. IEEE, 2008.
- [MGW10] Yu Ma, Xiaodong Gu, and Yuanyuan Wang. Histogram similarity measure using variable bin size distance. *Computer Vision and Image Understanding*, 114(8):981–989, August 2010.
- [MLB16a] J. Meyer, T. Längle, and J. Beyerer. Acquiring and processing light deflection maps for ransparent object inspection. In *2016 2nd International Conference on Frontiers of Signal Processing (ICFSP)*, pages 104–109, Oct 2016.
- [MLB16b] Johannes Meyer, Thomas Längle, and Jürgen Beyerer. About the acquisition and processing of ray deflection histograms for transparent object inspection. In *IRISH MACHINE VISION & IMAGE PROCESSING Conference proceedings*, 2016.
- [MS08] Kurt Mehlhorn and Peter Sanders. *Algorithms and data structures: The basic toolbox*. Springer Science & Business Media, 2008.
- [PHB13] Alexey Pak, Marco F. Huber, and Andrey Belkin. On weak distance between distributions in application to tracking. In *Sensor Data Fusion: Trends, Solutions, Applications (SDF), 2013 Workshop on*, pages 1–6. IEEE, 2013.
- [Sou14] César Souza. The Accord.NET Framework. <http://accord-framework.net>, 2014. Accessed: March 2017.
- [VJ04] Paul Viola and Michael J. Jones. Robust real-time face detection. *International journal of computer vision*, 57(2):137–154, 2004.