

Uncertainty Propagation using Copulas in a 3D Stereo Matching Pipeline

Roman Malinowski

Centre National d'Etudes Spatiales, Toulouse; CS Group, Toulouse ; UTC, Compiègne, France

ROMAN.MALINOWSKI@UTC.FR

Sébastien Destercke

Univeristé de Technologie de Compiègne, France

SEBASTIEN.DESTERCKE@UTC.FR

Emmanuel Dubois

Centre National d'Etudes Spatiales, Toulouse, France

EMMANUEL.DUBOIS@CNES.FR

Loïc Dumas

CS Group, Toulouse, France

LOIC.DUMAS@CSGROUP.EU

Emmanuelle Sarrazin

Centre National d'Etudes Spatiales, Toulouse, France

EMMANUELLE.SARRAZIN@CNES.FR

Abstract

This contribution presents a concrete example of uncertainty propagation in a stereo matching pipeline. It considers the problem of matching pixels between pairs of images whose radiometry is uncertain and modeled by possibility distributions. Copulas serve as dependency models between variables and are used to propagate the imprecise models. The propagation steps are detailed in the simple case of the Sum of Absolute Difference cost function for didactic purposes. The method results in an imprecise matching cost curve. To reduce computation time, a sufficient condition for conserving possibility distributions after the propagation is also presented. Finally, results are compared with Monte Carlo simulations, indicating that the method produces envelopes capable of correctly estimating the matching cost.

Keywords: imprecise probabilities, possibility distribution, copulas, uncertainty propagation, stereo matching

1. Introduction

This contribution presents a concrete example of uncertainty propagation in the context of photogrammetry, and more specifically in the crucial step of matching cost computation (Scharstein et al., 2001). Recent research in this field aims to estimate the uncertainty associated with dense stereo matching in specific steps of the pipeline (Xiaoyan Hu and Mordohai, 2012; Mehlretter and Heipke, 2021; Sarrazin et al., 2021), or with end-to-end methods (Mehlretter, 2020). Those methods either estimate the uncertainty *a posteriori* without considering the uncertainty of the input data, or are not explainable, as for the case of deep-learning based methods. Okutomi and Kanade (1994) have estimated the uncertainty using a precise density function, and adapt their matching strategy to minimize this uncertainty. Instead, we

argue in favor of using imprecise models to represent the uncertainty regarding our data, due to the noise and various processing steps that can affect the images. In this contribution, we will use belief functions, and more specifically possibility distributions (Dubois and Prade, 1992), to model the uncertainty on image intensities. Copulas are used to characterize the dependency between models of uncertainty, and will serve in the context of stereo matching to propagate the models through a cost function. The resulting belief function allows to define multiple envelopes with different degrees of plausibility, centered on the matching cost computed without uncertainty. Those envelopes are validated using Monte Carlo simulations with multiple models of noise on the input images.

Section 2 presents the stereo matching problem and the considered sources of uncertainty. Section 3 contains definitions regarding imprecise probabilities, dependency models and details the uncertainty model used on the input images. Section 4 explains the method used to propagate the imprecise models based on copulas and details the specific case of propagation through a matching cost function. A sufficient condition for conserving possibility distributions through the propagation is also described. Finally, we present the resulting belief functions as well as Monte Carlo simulations for different copulas in Section 5.

2. The Stereo Matching Problem

2.1. Dense Matching

Stereo Matching is one of the main steps when reconstructing 3D models from pairs of images by photogrammetry. For a complete description and review of stereo methods, we refer to (Scharstein et al., 2001). In photogrammetry, images of the same scene are acquired from different points




Figure 1: Homologous pixels in a pair of images

of view, and depth information of each pixel is retrieved by evaluating its displacement between images. Figure 1 presents a pair of images and highlights the position of the same pixel in both images. This displacement is called “disparity”. We consider the case with two images, referred to as the *left* and *right* images. Images are often rectified so that the displacement of pixels can only occur in one direction, usually horizontally (Fusiello et al., 2000). This allows to restrict the search for a pixel’s match to a single row. The problem could be briefly summarised as follows: given a pair of left-right images (I_L, I_R) , determine for every pixel $p_L(x, y) \in I_L$ of the left image the disparity d of position, allowing to find its homologous pixel $p_R(x, y - d) \in I_R$ in the right image. By knowing the disparity d of an object, the displacement B and focal length f of the camera, the depth z of the object is computed using the following relation:

$$z = \frac{Bf}{d} \quad (1)$$

In practice, not all pixels of the left image have a corresponding pixel in the right image. Some zones can become occluded by an object when changing the camera’s point of view, or similarly, pixels that were hidden behind an object in the left image might appear in the right image. Those occluded zones can be identified *a posteriori* (Fua, 1991), so we will not consider this issue here.

2.2. Cost Functions

For each pixel, we evaluate whether pixels of the other image are good candidates for a match by measuring their similarity. Cost functions measure the similarity between two patches centered on the pixels to be compared. Patches are usually squared windows, although other shapes can be considered (Buades and Facciolo, 2015). Evaluating the matching cost of a function between a patch in the left




Figure 2: Patch comparison to find the disparity between two images. From top to bottom: left image patch, patches from the right image, cost curve

image and a list of potential matches in the right image, gives a *cost curve*. Low values of the cost function represent a strong similarity, and the correct disparity is determined by finding the minimum of a cost curve. An example of this procedure is shown in Figure 2. The top and middle figures present a patch of the left image and a row of the right image, where a potential match is sought. The bottom plot shows the corresponding cost curve, where each patch of the right image row is compared with the left image patch. The minimum of the cost curve indicates the correct disparity.

In this article, we will focus on using a basic Sum of Absolute Differences (SAD) cost function, defined as follows. Given patches $W_L \subset I_L$ and $W_R \subset I_R$ of the same shape with n pixels (usually squares):

$$\text{SAD}(W_L, W_R) = \sum_{(p_i, q_i) \in (W_L, W_R)} |I_L(p_i) - I_R(q_i)| \quad (2)$$

where p_i and q_i are pixels at the same position i in their patch. For convenience purposes, we will refer to the Absolute Difference as AD. This cost function is not ideal, but it is preferred here for its didactic properties. An illustration of the SAD cost function can be found in Figure 3.

An ideal cost function would generate a cost curve with a unique minimum corresponding to the correct disparity. In practice, such a function is hard to determine. There is no guarantee that the minimum is unique, nor that it corresponds to the correct disparity. Different cost functions have been proposed to better identify the correct match




Figure 3: Example of the SAD cost function

(Hannah, 1994), some being robust to small variations of intensities (Zabih and Woodfill, 1994), or determined using advanced method such as deep learning approaches (Žbontar and LeCun, 2016). We have chosen the SAD method to focus on simplicity and to ease didactic explanations regarding uncertainty propagation. Other cost functions can produce better results but with less facility to explain this paper’s methodology.

2.3. Sources and Modeling of Uncertainty in Stereo Matching

In Section 2.1, we stated that images can be rectified to ensure horizontal displacement of pixels. It is often the case in remote-sensing, where images are taken by airplane or satellites (Michel et al., 2020). This pre-processing step, added to the noise of the sensor taking the image, generates uncertainty surrounding the value of every pixel. Our aim is to propagate this uncertainty through the matching cost evaluation, in order to have an estimation of the uncertainty attached to the matching cost of two patches. Knowing this uncertainty can help in better identifying the correct disparity (Okutomi and Kanade, 1994), for instance in the case where multiple minima of the cost curve exist, a good strategy may be to select the disparity with the least uncertainty. The following section presents models of uncertainty used in this contribution.

3. Uncertainty Models

The noise of the sensor and the uncertainty due to the pre-processing steps make it difficult to specify a precise probability model. Consequently, an *imprecise* model is preferred for this problem. Definitions regarding the imprecise probability framework are now presented.

3.1. Belief Functions

Consider a random variable X defined over a measurable space \mathcal{X} . A probability mass function m over a space \mathcal{X} is a mapping $\mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$ satisfying:

$$m(\emptyset) = 0, \quad \sum_{X \subseteq \mathcal{X}} m(X) = 1 \quad (3)$$

where $\mathcal{P}(\mathcal{X})$ is the power set of \mathcal{X} . The subsets of \mathcal{X} whose mass is strictly positive are referred with the letter a . For clarity, the exponent of focal sets will refer to the space or variable they are defined over, and the subscripts will refer to an order (if it exists). For instance, a_3^X refers to the third focal set of variable X . From a mass function, it is possible to define a belief Bel and a plausibility Pl function, which are mappings $\mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$ defined as:

$$\begin{aligned} \forall A \subseteq \mathcal{X}, \text{Bel}(A) &= \sum_{a \subseteq A} m(a) \\ \forall A \subseteq \mathcal{X}, \text{Pl}(A) &= \sum_{a, a \cap A \neq \emptyset} m(a) \end{aligned} \quad (4)$$

Similarly, a belief function can also be interpreted as the lower bound of a convex set of probability measures \mathcal{M} , called a *credal set*, defined as:

$$\mathcal{M} = \{P \mid P(A) \geq \text{Bel}(A), \forall A \subseteq \mathcal{X}\} \quad (5)$$

3.2. Possibility Distributions on Image Intensities

We work with grayscale images with a quantification of intensity levels in $[0, 255]$, corresponding to our measurable space \mathcal{X} . In this paper, we make the hypothesis that the value of a pixel cannot be more than 1 intensity level away from its observed value, and that the observed value is the most plausible. This is due to the quantification of the observed radiometry into integers. We thus chose to model the uncertainty of the intensity of every pixel $p \in I_L, I_R$ by a possibility distribution π centered on the observed intensity i_p :

$$\pi(i_p) = 1, \quad \pi(i_p \pm 1) = \alpha \quad (6)$$

with $\alpha \in [0, 1]$. In our simulation, $\alpha = 0.33$ for pixels in the left image, and $\alpha = 0.4$ for pixels in the right image. We use different values of *alpha* for the left and right image because the uncertainty model could change from

one image to the other, due to a different exposure, different noise or different calibration of our camera. This model is equivalent to state that we accept every probability with support in $[i_p - 1, i_p + 1]$ and whose probability measure P verifies $\{P(A) \leq \sup_{i \in A} \pi(i)\}$ as an acceptable model for our uncertainty. It has been shown (Dubois and Prade, 1992) that to every possibility distribution corresponds a minitive belief function (also called a necessity function) whose mass distribution function for every focal set a^P is in our case:

$$\begin{aligned} m_p(a_1^P = \llbracket i_p, i_p \rrbracket) &= 1 - \alpha \\ m_p(a_2^P = \llbracket i_p - 1, i_p + 1 \rrbracket) &= \alpha \end{aligned} \quad (7)$$

with $\llbracket \cdot, \cdot \rrbracket$ referring to integer intervals. In particular, $\llbracket i_p, i_p \rrbracket$ correspond to the singleton $\{i_p\}$

It is noteworthy that in this problem of disparity estimation, we only consider the uncertainty of our input image intensities, but do not consider the uncertainty regarding our cost function's ability to correctly identify the true disparity as its minimum. In other words, we don't take into consideration the uncertainty resulting from the difference between "two patches are very similar" and "the pixels at the center of the patches are homologous". To better illustrate this, consider a case where two pixels should be matched, but the pixels in the patches surrounding them are dissimilar. Then the cost function between those two patches would be high, and there can be another patch with a lower cost function that would be wrongly selected as it is the minimum of the cost curve.

3.3. Copulas as Dependency Models

When propagating probability densities, one has to take into account the dependency between the different sources of uncertainty. Copulas are great tools to model the dependency between variables, as it has been shown that they can represent any kind of dependency (Sklar, 1959). A copula, or N -copula, is a mapping $C : [0, 1]^N \rightarrow [0, 1]$ satisfying a number of properties (Nelsen, 2006) expressed here in the N dimensional case. For all k in $\llbracket 1, N \rrbracket$ and for all (u_1, \dots, u_N) in $[0, 1]^N$:

$$C(u_1, \dots, u_{k-1}, 0, u_{k+1}, \dots, u_N) = 0 \quad (8)$$

$$C(1, \dots, 1, u_k, 1, \dots, 1) = u_k \quad (9)$$

It also needs to be N -increasing, meaning that for all $U = (u_1, \dots, u_N)$ in $[0, 1]^N$, and for all $V = (v_1, \dots, v_N)$ in $[0, 1]^N$, such that $\forall k \in \llbracket 1, N \rrbracket, u_k \leq v_k$:

$$\sum_{\substack{(w_1, \dots, w_N) \in \\ \otimes_k \{u_k, v_k\}}} (-1)^{|\{k \mid w_k = u_k\}|} C(w_1, \dots, w_k) \geq 0 \quad (10)$$

with \otimes representing the Cartesian product of sets, and $|\cdot|$ corresponding to the cardinal of a set. The left term of

the previous inequality is also called the hyper-volume or H -volume of the copula. The inequality then reads as the H -volume of the copula of every segment of the unit hypercube is positive.

A copula can also be interpreted as a multivariate distribution function whose marginals are uniform on the unit interval.

Some famous copulas include:

- The product copula representing independence: $C_\Pi = \prod_k u_k$
- The upper Fréchet-Hoeffding bound representing complete co-monotonicity: $C_{\min} = \min_k u_k$
- the Gaussian copula with correlation matrix R : $C_R = \Phi_R(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_N))$, where Φ_R is the joint multivariate distribution function of a Gaussian variable with correlation matrix R , and Φ^{-1} is the inverse distribution function of a univariate Gaussian variable.

Sklar's theorem states that every multivariate cumulative distribution function (CDF) $G : \otimes_k \mathcal{X}_k \rightarrow [0, 1]$ can be expressed by means of its marginals CDF $F_k : \mathcal{X}_k \rightarrow [0, 1], k \in \llbracket 1, N \rrbracket$ and a copula C :

$$G = C(F_1, \dots, F_N) \quad (11)$$

The reverse implication is also true, meaning that joining any univariate CDFs with a copula returns a correctly define multivariate CDF.

3.4. Sampling from Copulas

In our case, we want to model with a copula the dependency between the random intensities of two pixels: one in the left image, and one in the right image. We propose to model their dependency with the product copula if the pixels are not from the same physical object (meaning that the value of their intensities are independent), and by a Gaussian copula with a covariance matrix $\begin{pmatrix} 1 & \sigma_{\text{obj}} \\ \sigma_{\text{obj}} & 1 \end{pmatrix}$, $\sigma_{\text{obj}} \in \mathbb{R}^+$, if they belong to the same physical object in the scene. A segmentation of the image based on the ground truth of the disparity is used to determine if two pixels belong to the same object. As we validate our method using Monte Carlo simulations in Section 5, we need to be able to sample random vectors from a copula. We now detail a method for sampling from copulas in general and a method for sampling from the Gaussian copula. Given a copula C , and two cumulative distribution functions F_X and F_Y , the method used to generate a pair of observations (x, y) from a joint CDF $C(F_X, F_Y)$ is the following:

- Sample two independent samples u_1, u_2 from a uniform distribution on $[0, 1]$

- Set $v = \partial C^{-1}(u_2)$ where ∂C^{-1} is the quasi-inverse of the partial derivative of C regarding its first variable (which exists almost everywhere and is invertible).
- (u_1, v) each follow a uniform distribution on $[0, 1]$, and their associated copula is C
- The desired pair is $(x_1, x_2) = (F_X^{-1}(u_1), F_Y^{-1}(v))$, with F_X^{-1}, F_Y^{-1} being the quasi-inverses of the marginals CDFs.

Details of this method in the N -dimensional case can be found in (Cherubini et al., 2004). Simulation draws from the Gaussian N -copula with correlation matrix R are simpler to obtain:

- Compute the Cholesky decomposition A of the correlation matrix R
- Draw N independent random samples $u = (u_1, \dots, u_N)'$ from $\mathcal{N}(0, 1)$
- Set $v = Au$
- Set $w_k = \Phi(v_k)$ where Φ is the univariate normal distribution function
- The desired draw is $(x_1, \dots, x_N) = (F_1^{-1}(w_1), \dots, F_N^{-1}(w_N))$ with F_k^{-1} , being the quasi-inverse of the k -th marginal CDF.

4. Using Copulas to Propagate Uncertainty

4.1. Combining Belief Functions with Copulas

In the following paragraph, we will explain how we use copulas to propagate the uncertainty. Let us first explain it in the precise case. Let $\mathcal{X} = x_1, \dots, x_{N^X}$, $\mathcal{Y} = y_1, \dots, y_{N^Y}$ and \mathcal{Z} be three discrete spaces, and let X, Y be two discrete random variables taking values in \mathcal{X}, \mathcal{Y} respectively, with respective CDFs $F_X : \mathcal{X} \rightarrow [0, 1]$, $F_Y : \mathcal{Y} \rightarrow [0, 1]$ and whose dependency can be represented by a copula C . To avoid heavy notations, we will refer to $\Delta C_{u_1, v_1}^{u_2, v_2}$ as the H -volume of the copula C over $[u_1, u_2] \otimes [v_1, v_2] \subseteq [0, 1]^2$. Then we know that the joint probability mass distribution $p : \mathcal{X} \otimes \mathcal{Y} \rightarrow [0, 1]$ is defined as

$$\begin{aligned} \forall (i, j) \in \llbracket 1, N^X \rrbracket \otimes \llbracket 1, N^Y \rrbracket, \\ p(x_i, y_j) = \Delta C_{F_X(x_{i-1}), F_Y(y_{j-1})}^{F_X(x_i), F_Y(y_j)} \end{aligned} \quad (12)$$

with the convention that $F_X(x_0) = F_Y(y_0) = 0$.

Let $f : \mathcal{X} \otimes \mathcal{Y} \rightarrow \mathcal{Z}$ be a mapping and we define the random variable Z as $Z = f(X, Y)$. Then the probability mass distribution p_Z of Z is:

$$\forall z \in \mathcal{Z}, p_Z(z) = \sum_{\substack{x, y \\ z=f(x, y)}} p(x, y) \quad (13)$$

Determining every (x, y) , whose image by f equals z , is not always trivial. This becomes even more complex when we are considering copulas with $N > 2$ variables. Note that the H -volume is a sum of 2^N terms, which also increases exponentially with the dimension. In the continuous case, the H -volume is replaced with the density h of the joint CDF, and the density of Z is

$$p_Z(z) = \int_{\mathcal{X}} \int_{\mathcal{Y}} h(x, y) \text{Ind}(f(x, y) = z) dx dy \quad (14)$$

where Ind is an indicator function.

Let us now consider the case where we do not know the CDFs of X and Y , but only the belief functions $\text{Bel}_X, \text{Bel}_Y$ and their associated mass functions $m_X : \mathcal{P}(\mathcal{X}) \rightarrow [0, 1]$, $m_Y : \mathcal{P}(\mathcal{Y}) \rightarrow [0, 1]$. Propagating belief functions can be done in similar manner as in (12), replacing the CDFs by cumulated masses (Ferson et al., 2004). In order to do so, we need to fix an arbitrary order over the focal sets of m_X and m_Y . Let us suppose we have such an order and that the focal sets of m_X are $\{a_1^X, \dots, a_{N_a^X}^X\}$ and those of m_Y are $\{a_1^Y, \dots, a_{N_a^Y}^Y\}$. Then the joint mass for X and Y is defined as:

$$m_{XY}(a_i^X, a_j^Y) = \Delta C_{\sum_{k=1}^{i-1} m_X(a_k^X), \sum_{k=1}^{j-1} m_Y(a_k^Y)}^{\sum_{k=1}^i m_X(a_k^X), \sum_{k=1}^j m_Y(a_k^Y)} \quad (15)$$

We refer to a previous publication for details about this method and details on other ways of aggregating credal sets with a copula, as well as the importance of the order over focal sets (Malinowski and Destercke, 2023).

An important thing to note here is that a copula does not bare the same meaning when used with precise models in (12) and when used with imprecise models as in (15). In the precise case, the copula will encode the dependency between the values of the random variables. For instance in the case of co-monotonicity inside an image patch, the minimum copula C_{\min} would mean that a high value of a pixel indicates that the pixel's neighbours must have also possess high intensities. In the imprecise case, the copula will encode the dependency between degrees of belief regarding the values of random variables (supposing that there exist such underlying random variables behind our models). For instance the minimum copula C_{\min} indicates that a high belief regarding a pixel's intensity is correlated to a high belief regarding its neighbours. However, the (confident) values of pixels could be very low and very high, which would not be the case using (12).

Given the joint mass distribution function constructed with the copula, it is then possible to compute the mass distribution function of Z :

$$\forall a^Z \subseteq \mathcal{Z}, m_Z(a^Z) = \sum_{\substack{a_i^X, a_j^Y \\ a^Z=f(a_i^X, a_j^Y)}} m_{XY}(a_i^X, a_j^Y) \quad (16)$$

As in the precise case, computing the image of f for every pair of focal sets (a_i^X, a_j^Y) is not always trivial.

4.2. Uncertainty Propagation through the SAD Cost Function

To illustrate how to propagate the uncertainty using belief functions and a copula, we will present in this section the case of the uncertainty related to the SAD cost function.

The SAD is computed between two 3×3 windows W_L, W_R . The propagation of uncertainty is done in two passes: in the first pass, the pixel-to-pixel Absolute Difference (AD) of intensities is computed, resulting in a single 3×3 window. In the second pass, all of the AD are summed, resulting in the complete SAD. Figure 3 illustrates those two steps. For the AD, the uncertainty for each ‘‘pixel’’ is computed using a Gaussian 2-copula characterizing the dependency between each pair of pixels. We use the mass distribution m_p of Equation (7) to represent the uncertainty of each pixel p . For every pair of pixel $p \in I_L, q \in I_R$, we note $AD_{pq} = |i_p - i_q|$. There exists 3 focal sets related to the AD:

- a_1^{AD} is obtained by computing the AD of a_1^p and a_1^q
- a_2^{AD} is obtained by computing the AD of a_2^p and a_1^q or a_1^p and a_2^q
- a_3^{AD} is obtained by computing the AD of a_2^p and a_2^q

To compute their exact image through the AD, we need to take into account the non monotonicity of the absolute value around 0:

$$\begin{aligned} a_1^{\text{AD}} &= [[AD_{pq}, AD_{pq}]] \\ a_2^{\text{AD}} &= [[AD_{pq} - 1, AD_{pq} + 1]] \text{ if } AD_{pq} > 0 \\ &= [[AD_{pq}, AD_{pq} + 1]] \text{ otherwise} \\ a_3^{\text{AD}} &= [[AD_{pq} - 2, AD_{pq} + 2]] \text{ if } AD_{pq} > 1 \\ &= [[AD_{pq} - 1, AD_{pq} + 2]] \text{ if } AD_{pq} = 1 \\ &= [[AD_{pq}, AD_{pq} + 2]] \text{ otherwise} \end{aligned}$$

The mass of each focal set of AD is computed using Equations (15) and (16), and by remarking that $m_p(a_1^p) + m_p(a_2^p) = 1$ and $m_q(a_1^q) + m_q(a_2^q) = 1$:

$$\begin{aligned} m_{\text{AD}}(a_0^{\text{AD}}) &= \Delta C_{0,0}^{m_p(a_1^p), m_q(a_1^q)} \\ m_{\text{AD}}(a_1^{\text{AD}}) &= \Delta C_{m_p(a_1^p), m_q(a_1^q)}^{1, m_q(a_1^q)} + \Delta C_{m_p(a_1^p), m_q(a_1^q)}^{m_p(a_1^p), 1} \\ m_{\text{AD}}(a_2^{\text{AD}}) &= \Delta C_{m_p(a_1^p), m_q(a_1^q)}^{1, 1} \end{aligned}$$

Once all of the AD focal sets and mass functions are computed, they are summed over the 3×3 window. The same steps are repeated, this time with 9 sources of uncertainty. The dependency between all of the pixels is represented by a Gaussian 9-copula.

4.3. Reducing the Computation Time of Exact Propagation

Determining the images of the AD focal sets by the sum is trivial as we only need to sum the bounds of each interval for every combination of them. Computing the joint mass over a 3×3 window is not as easy as in the AD case. The H -volume is now computed for a 9-copula, which is a sum of 2^9 terms. Because the uncertainty of each absolute difference is represented by 3 focal sets, there are 3^9 combinations of focal sets to evaluate for the whole 3×3 window. A way of reducing the computation time is to take advantage of the fact that the focal sets derived from a possibility distribution (or equivalently from its necessity measure) form a nested family of sets. In the general case, propagating two necessity measures $\text{Nec}_X : \mathcal{P}(X) \rightarrow [0, 1], \text{Nec}_Y : \mathcal{P}(Y) \rightarrow [0, 1]$ through a mapping $f : X \times Y \rightarrow Z$ with a copula C does not yield a necessity measure but ‘‘only’’ a belief function. For the special case where $\text{Nec}_X, \text{Nec}_Y$ are defined by symmetric uni-modal possibility distributions (typically triangular possibilities), and f is a monotone function applied to a linear combination $\alpha X + \beta Y + \gamma$ of X and Y , $(\alpha, \beta, \gamma) \in \mathbb{R}^3$, then the focal sets of $Z = f(X, Y)$ form a nested family of sets, which is characteristic of necessity measures (Shafer, 1976). Indeed, the focal sets of $\text{Nec}_X, \text{Nec}_Y$ are families of nested sets that can be written as $([\bar{X} \pm \Delta x_i]), ([\bar{Y} \pm \Delta y_j])$, with $\bar{X} \in X, \bar{Y} \in Y$ and $(\Delta x_i), (\Delta y_j)$ positive scalars. The focal sets of the linear combination of $\alpha X + \beta Y + \gamma$ are of the form:

$$\begin{aligned} a_{ij} &= [\alpha \bar{X} + \beta \bar{Y} + \gamma - (|\alpha| \Delta x_i + |\beta| \Delta y_j), \\ &\quad \alpha \bar{X} + \beta \bar{Y} + \gamma + (|\alpha| \Delta x_i + |\beta| \Delta y_j)] \end{aligned}$$

and is a nested family of sets. Applying a monotone function to those focal sets will keep the nesting property (but not their symmetry), which results in Bel_Z being a necessity measure. For more advanced functions such as multiplication, exponential etc. . . , this property is not always true (it is easy to find examples where the nesting property is not retained for Bel_Z). This property can be used to simplify and reduce the computations of focal sets bounds and their masses in the case where the AD is superior to 2 (to avoid the non monotonicity of the absolute value around 0). In the case where the focal sets are nested, it holds that for every focal sets a_i^X, a_j^Y of $\text{Nec}_X, \text{Nec}_Y$ (Malinowski and Destercke, 2023):

$$\text{Bel}_{XY}(a_i^X, a_j^Y) = C(\text{Nec}_X(a_i^X), \text{Nec}_Y(a_j^Y)) \quad (17)$$

hence the joint belief function can be computed solely with the marginal masses of Nec_X and Nec_Y and the copula. There is no need to compute the joint mass, thus avoiding the computation of the H -volume with (10).

5. Resulting Envelopes and Monte Carlo Simulations

We have presented the tools used for propagating belief functions through the cost function. This section details the construction of the correlation matrix for the Gaussian copula, and how envelopes were computed from the propagated belief function. We also generate samples from different noise models with a dependency specified by the product and Gaussian copulas. This noise is added to the input images, and we apply the Monte Carlo method to generate multiple “noised” cost curves. Those curves are compared to the envelopes to determine if the propagated belief functions is able to correctly characterize the uncertainty.

5.1. Constructing a Correlation Matrix for the Gaussian Copula

Section 3.4 presented our aim to sample from copulas to validate the computed envelopes. The current section details the construction of the Gaussian copula which will be used. The Gaussian copula is parameterized by a correlation matrix R . The correlation between the uncertain sources is based on a segmentation $S : (I_L \cup I_R) \rightarrow [[0, K]]$, $K \in \mathbb{N}$, of the images, performed on the ground truth disparity of each pixel. Given two pixels $(p, q) \in (I_L \cup I_R)^2$, their covariance is determined by:

$$\text{cov}(p, q) = \begin{cases} 1 & \text{if } p = q \\ \rho_k, & \text{if } p \neq q \text{ and } S(p) = S(q) \\ 0, & \text{otherwise} \end{cases} \quad (18)$$

The segmentation contains $K = 8$ different clusters, and for every k in $[[0, K]]$, ρ_k is assigned an arbitrary value between 0.9 and 1.

We can thus build a correlation matrix R for all pixels of $(I_L \cup I_R)$, and sample a perturbation on all of the pixels using a Gaussian copula with this correlation matrix. The Gaussian 2-copula and Gaussian 9-copula are used to model the dependency between masses in the uncertainty propagation step and their correlation matrix is also constructed using Equation (18).

5.2. Envelopes from the Propagated Belief Function

To illustrate the propagation of uncertainty with a concrete example, we used the *Middlebury* dataset¹, consisting of pairs of left-right images with the correct disparity map available. An example of an image pair from this dataset is presented in Figure 1. For every pixel in the left image, we computed its SAD cost curve, while propagating the uncertainty model presented in (3.2). The focal sets representing

the uncertainty related to the SAD value at every considered disparity are intervals containing the “precise” SAD value. Upper and lower envelopes with different plausibility levels have been computed. They represent the biggest (resp. lowest) focal set bound \bar{a} whose plausibility, computed using Equation (4), is above a given threshold T : $\text{Pl}(\bar{a}) \geq T$. The plausibility threshold 0 is strict $\text{Pl}(\bar{a}) > 0$ and represents the support of the SAD. The plausibility threshold 1 coincides with the value of the cost curve obtained in the “precise setting” (i.e. without considering the uncertainty).

Figure 4 contains the cost curve and its support obtained using the product copula. The curves obtained using the Gaussian copula with correlation matrix R present the exact same values, so they are not represented here. In this case; this is due to the fact that the Gaussian Copula C_R and the product copula C_{Π} never assign a null value to a joint mass. Thus the joint belief functions computed using those copulas have the same focal sets. Taking the product copula and, for example, the lower Fréchet-Hoeffding 2-copula $C(u, v) = \max(u + v - 1, 0)$, might have led to different joint focal sets, and thus to different envelopes.


In the specific case of the product copula C_{Π} , the joint mass is easy to compute as it is simply the product of the marginal masses. The black dotted rectangle in Figure 4(a) represents a region of interest (ROI), defining the bounds of Figures 5 and 6. The green vertical line represents the correct disparity. It is important to remark that in the “precise” case, the minimum of the SAD curve does not indicate the true disparity as there is a slight offset between this minimum and the ground truth. When considering the support of the cost curve (where $\text{Pl} > 0$), it is possible to determine a set of disparities containing the true value of the disparity. By noting $\underline{\text{SAD}}(d)$ (resp. $\overline{\text{SAD}}(d)$) the upper (resp. lower) envelope of the cost curve at disparity d , the set containing the true disparity is estimated as:

$$D = \{d \mid \underline{\text{SAD}}(d) \leq \min_{d'} \overline{\text{SAD}}(d')\} \quad (19)$$

This set is represented graphically by the blue arrow in the Figure 4(b). This set can later be used to determine confidence intervals for the true disparity (or for the true depth using eq. (1)), or could be used to better estimate the true disparity by solving an optimization problem under uncertainty. We do not cover these topics in this article.


Figure 5 displays more detailed plausibility thresholds. The envelopes of Figure 5(a) where computed using the product copula C_{Π} , whereas those of Figure 5(b) where computed using the Gaussian copula C_R . As stated previously, the envelopes corresponding to plausibility levels 0 and 1 are the same for both copulas. However, the plausibility levels 0.99 and 0.6 are more concentrated around plausibility level 1 in the case of the Gaussian copula than in the case of the product copula. By construction, the Gaussian copula

¹<https://vision.middlebury.edu/stereo/data/scenes2003/>




(b) Focus around the true disparity and the possible disparity set D

Figure 4: Cost curve ($PI = 1$) and its support ($PI > 0$) using the product copula C_{Π} . The green vertical line represents the true disparity



(a) Different plausibility levels for the product copula C_{Π}



(b) Different plausibility levels for the Gaussian copula C_R (the $PI = 1$ and $PI \geq 0.99$ curves overlap)

Figure 5: Zoom over a ROI and plausibility levels

C_R is more comonotone than the product copula, which could explain this effect.

5.3. Comparison with Monte Carlo Samplings

Figure 6(a) contains Monte Carlo samplings for independent centered normal noise (with standard deviation $\sigma = 1$) and independent uniform noise (with support $[-1, 1]$). Figure 6(b) contains the same type of noise models, but this time sampled from the Gaussian copula as their dependency model as in Section 3.4. Different plausibility levels computed with their respective copula are also plotted. We observe that all Monte Carlo samplings, no matter the noise model used, are correctly contained in the support envelopes. This holds for both the product and the Gaussian copulas. It suggests that the method used for computing the envelopes allows to correctly estimate the possible values of the cost curve when the input images are uncertain. The envelopes computed with the product copula are pessimist, as the independent Monte Carlo samplings are rarely outside the 0.6 plausibility level. The envelopes obtained with the Gaus-

sian copula suffer less from this observation, with more samples venturing outside the envelopes. Nevertheless, the method for propagating belief functions correctly provides confidence levels regarding the value of the cost function at all the disparities.

6. Conclusion


This contributions presents a real-life application of uncertainty propagation using possibility distributions as models, and copulas to characterize the dependency between different random sources. In order to propagate the uncertainty in the matching step of a photogrammetry 3D pipeline, we introduce the use of cost functions, and presented a simple model to represent the sources of uncertainty in input images. The different steps for propagating the uncertainty are detailed for didactic purposes, with the intention of highlight the potential of using imprecise models in concrete cases. Additionally, a sufficient condition for conserving possibilities after the propagation is proposed, which can be used to reduce the computation time.

Envelopes are deduced from the propagated plausibility functions, which correctly frame different types of input noise, generated using Monte Carlo simulations. Although copulas do not bear the same meaning when used with precise density functions and when used with belief functions, simulations show that the propagated belief functions allow to generate correct envelopes which estimate the possible value of the cost function. Dealing with uncertain cost functions allows as well to estimate a set of acceptable disparities, which correctly contains the true disparity.


Although we presented a case of using imprecise probability to estimate the cost function in a photogrammetry 3D pipeline problem, we did not consider the uncertainty stemming from the formulation of the function itself. It is not guaranteed that comparing two patches of images as we did is sufficient to determine with certainty all disparities. Future work includes taking into account both the uncertainty on the input image and the uncertainty related to the cost function's ability to correctly distinguish homologous pixels. Another perspective is to extend this work to more complex cost functions and to other step of the 3D pipeline, such as rasterisation of imprecise 3D point clouds or stereo-rectification of pairs of images.

Acknowledgments

This project has received financial support from the CNRS through the MITI interdisciplinary program and from CNES. CNES and CS Group also finance a PhD from which this work is derived.



(a) Envelopes and Monte Carlo sampling using the product copula C_{Π}



(b) Envelopes and Monte Carlo sampling using the Gaussian copula C_R

Figure 6: Zoom over a ROI and Monte Carlo sampling

Author Contributions

Roman Malinowski has worked on the methodology, done the implementation and simulations of the method presented in this contribution, as well as most of the writing. Sébastien Destercke helped for the writing as well as the methodology regarding uncertainty modelling and propagation. Emmanuel Dubois, Loïc Dumas and Emmanuelle Sarrazin helped during the writing, especially regarding the photogrammetry aspects, and with the stereo Matching problems. Overall, every idea leading to this work emerged from numerous discussions between all the authors.

References

- Antoni Buades and Gabriele Facciolo. Reliable multiscale and multiwindow stereo matching. 8(2):888–915, 2015. ISSN 1936-4954. doi:[10.1137/140984269](https://doi.org/10.1137/140984269).
- Umberto Cherubini, Elisa Luciano, and Walter Vecchiato. *Simulation of Market Scenarios*, chapter 6, pages 181–194. John Wiley and Sons, Ltd, 2004. ISBN 9781118673331. doi:<https://doi.org/10.1002/9781118673331.ch6>.
- Didier Dubois and Henri Prade. When upper probabilities are possibility measures. 49(1):65–74, 1992. ISSN 01650114. doi:[10.1016/0165-0114\(92\)90110-P](https://doi.org/10.1016/0165-0114(92)90110-P).
- Scott Ferson, William Oberkampf, W Tucker, Jianzhong Zhang, Lev Ginzburg, Daniel Berleant, Janos Hajagos, and Roger Nelsen. Dependence in probabilistic modeling, dempster-shafer theory, and probability bounds analysis., 2004.
- Pascal Fua. Combining stereo and monocular information to compute dense depth maps that preserve depth discontinuities. 1991.
- Andrea Fusiello, Emanuele Trucco, and Alessandro Verri. A compact algorithm for rectification of stereo pairs. *Machine Vision and Applications*, 12(1):16–22, July 2000. ISSN 0932-8092, 1432-1769. doi:[10.1007/s001380050120](https://doi.org/10.1007/s001380050120).
- Marsha Jo Hannah. Computer matching of areas in stereo images, 1994.
- Roman Malinowski and Sébastien Destercke. Copulas, lower probabilities and random sets: How and when to apply them? In Luis A. García-Escudero, Alfonso Gordaliza, Agustín Mayo, María Asunción Lubiano Gomez, Maria Angeles Gil, Przemyslaw Grzegorzewski, and Olgierd Hryniewicz, editors, *Building Bridges between Soft and Statistical Methodologies for Data Science*, pages 271–278. Springer International Publishing, 2023. ISBN 978-3-031-15509-3.
- Max Mehlretter. Uncertainty estimation for end-to-end learned dense stereo matching via probabilistic deep learning, 2020. Number: arXiv:2002.03663.
- Max Mehlretter and Christian Heipke. Aleatoric uncertainty estimation for dense stereo matching via CNN-based cost volume analysis. 171:63–75, 2021. ISSN 09242716. doi:[10.1016/j.isprsjsprs.2020.11.003](https://doi.org/10.1016/j.isprsjsprs.2020.11.003).
- J. Michel, E. Sarrazin, D. Youssefi, M. Cournet, F. Buffe, J. M. Delvit, A. Emilien, J. Bosman, O. Melet, and C. L’Helguen. A new satellite imagery stereo pipeline designed for scalability, robustness and performance. V-2-2020:171–178, 2020. ISSN 2194-9050. doi:[10.5194/isprs-annals-V-2-2020-171-2020](https://doi.org/10.5194/isprs-annals-V-2-2020-171-2020).
- Roger B. Nelsen. *An introduction to copulas*. Springer series in statistics. Springer, 2. ed edition, 2006. ISBN 978-0-387-28659-4 978-1-4419-2109-3.
- Masatoshi Okutomi and Takeo Kanade. A stereo matching algorithm with an adaptive window: theory and experiment. 16(9):920–932, 1994. ISSN 01628828. doi:[10.1109/34.310690](https://doi.org/10.1109/34.310690).
- E. Sarrazin, M. Cournet, L. Dumas, V. Defonte, Q. Fardet, Y. Steux, N. Jimenez Diaz, E. Dubois, D. Youssefi, and F. Buffe. Ambiguity concept in stereo matching pipeline. XLIII-B2-2021:383–390, 2021. ISSN 2194-9034. doi:[10.5194/isprs-archives-XLIII-B2-2021-383-2021](https://doi.org/10.5194/isprs-archives-XLIII-B2-2021-383-2021).
- D. Scharstein, Richard Szeliski, and R. Zabih. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. In *Proceedings IEEE Workshop on Stereo and Multi-Baseline Vision (SMBV 2001)*, pages 131–140. IEEE Comput. Soc, 2001. ISBN 978-0-7695-1327-0. doi:[10.1109/SMBV.2001.988771](https://doi.org/10.1109/SMBV.2001.988771).
- Glenn Shafer. *A Mathematical Theory of Evidence*. 1976.
- M. Sklar. *Fonctions de Répartition À N Dimensions Et Leurs Marges*. Université Paris 8, 1959.
- Xiaoyan Hu and P. Mordohai. A quantitative evaluation of confidence measures for stereo vision. 34(11):2121–2133, 2012. ISSN 0162-8828, 2160-9292. doi:[10.1109/TPAMI.2012.46](https://doi.org/10.1109/TPAMI.2012.46).
- Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In Jan-Olof Eklundh, editor, *Computer Vision — ECCV ’94*, volume 801, pages 151–158. Springer Berlin Heidelberg, 1994. ISBN 978-3-540-57957-1 978-3-540-48400-4. doi:[10.1007/BFb0028345](https://doi.org/10.1007/BFb0028345). Series Title: Lecture Notes in Computer Science.

Jure Žbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. 2016.