Constructing a braid of partitions from hierarchies of partitions

Guillaume Tochon[1], Mauro Dalla Mura[2,3], and Jocelyn Chanussot[2]

[1] EPITA Research and Development Laboratory (LRDE), Le Kremlin-Bicêtre, France
guillaume.tochon@lrde.epita.fr
[2] GIPSA-lab, Grenoble Institute of Technology, Saint Martin d'Hères, France
{mauro.dalla-mura, jocelyn.chanussot}@gipsa-lab.grenoble-inp.fr
[3] Tokyo Tech World Research Hub Initiative (WRHI), School of Computing, Tokyo Institute of Technology, Tokyo, Japan

Abstract. Braids of partitions have been introduced in a theoretical framework as a generalization of hierarchies of partitions, but practical guidelines to derive such structures remained an open question. In a previous work, we proposed a methodology to build a braid of partitions by experimentally composing cuts extracted from two hierarchies of partitions, notably paving the way for the hierarchical representation of multimodal images. However, we did not provide the formal proof that our proposed methodology was yielding a braid structure. We remedy to this point in the present paper and give a brief insight on the structural properties of the resulting braid of partitions.

Keywords: braid of partitions, hierarchy of partitions, h-equivalence, multimodal images

1 Introduction

Hierarchical representations are a well suited tool to handle the multi-scale nature of images, since they allow to encompass all potential scales of interest in a single structure. The hierarchical representation can be constructued once and regardless of the application, and its scale of analysis can then be tuned afterward to comply with the pursued task. The component tree (also called min-tree and max-tree) [14], the inclusion tree (also called tree of shapes (ToS)) [7], the α-tree (also called the hierarchy of quasi-flat zones) [16] and the binary partition tree (BPT) [13] constitute a non-exhaustive list of the most known hierarchical representations in the mathematical morphology literature. Reviews can be found in [3, 8]. Hierarchical representations have proven to be useful for many image processing and computer vison tasks, such as image filtering [20] and simplification [21], image segmentation [11] as well as object recognition [19]. The most common framework is to build and process a single hierarchical representation for a given input image. In some cases however, it could be interesting to associate one image with multiple hierarchical representations (each one focusing on a particular feature of the image for example), or, on the contrary, to build

a common hierarchical representation for multiple input images (each being a single modality of a multimodal image for instance). While these largely remain open questions, some recent works have been devoted to such fusion issues. The fusion of multiple hierarchical representation is for instance addressed in [5], where hierarchies of watersheds (see [2]) driven by area and dynamics attributes are combined through the composition by infimum, supremum or averaging of their saliency maps. The representation of multimodal images (*i.e.* several images acquired over the same scene with different setups, such as different sensor types or acquisition dates) with a single hierarchical structure is another challenge studied in the literature [17]. The ToS structure has for instance been extended to multivariate images in [4], where univariate ToS are first computed for each individual modality and then further merged into a graph from which is derived the final multivariate ToS representation (note that a similar idea is presented in [10] to extend component trees to multivariate images, but the final result is a graph and no longer a tree-based representation). In [9], a single BPT is built over a whole video sequence by integrating motion cues during the construction stage, allowing to perform some object tracking by simply identifying nodes of interest in the resulting trajectory BPT structure. Another approach for the construction of a multi-feature BPT has been introduced in [12], where all modalities of the input multimodal image cooperate in a consensus framework to allow for the construction of a single tree structure. Finally, braids of partitions were proposed in [6] as a generalization of hierarchies of partitions for a theoretical standpoint, and we actually sketched in a previous work [18] the potentiality of such braid structures to act as suited hierarchical representations of multimodal images. In [18], we proposed to build the braid structure (and its associated monitor hierarchy) by experimentally combining cuts extracted from two hierarchies of partitions, and showed the interest of the resulting structure within the framework of multimodal image segmentation. However, we did not provide the formal proof that the proposed methodology yielded a braid structure.
We remedy to this point in this present article, whose organization is as follows: Section 2 introduces all used notations and formally defines the notions of hierarchies and braids of partitions. Section 3 recalls the construction procedure introduced in [18] and formally proves that the resulting structure is a braid. Section 4 gives a few insights on the structural properties of braids obtained with the presented construction procedure while Section 5 concludes and presents the perspectives of the actual work.

2 From hierarchies to braids of partitions

2.1 Hierarchies of partitions

Let E be the spatial support of a generic image, *i.e.*, its pixel grid (in which case $E \subseteq \mathbb{Z}^2$ although there is no requirement for E to be discrete in the following). A partition π of E is a collection of regions $\{\mathcal{R}_i \subseteq E\}$ of E such that $\mathcal{R}_i \cap \mathcal{R}_{j \neq i} = \emptyset$ and $\bigcup_i \mathcal{R}_i = E$. The set of all possible partitions of E is denoted Π_E. For any two partitions $\pi_i, \pi_j \in \Pi_E$, $\pi_i \leq \pi_j$ when for each region $\mathcal{R}_i \in \pi_i$, there exists

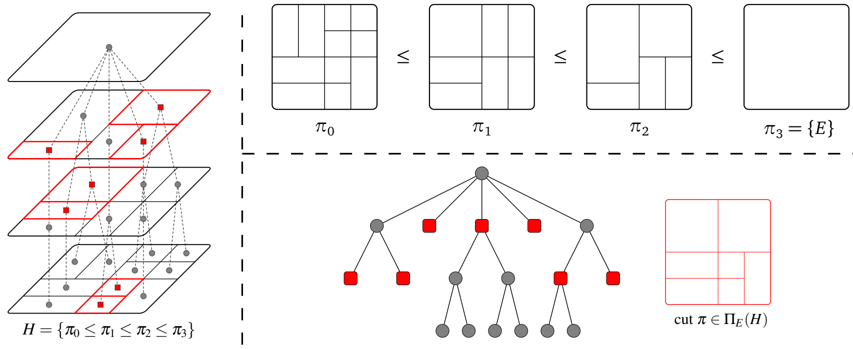

Fig. 1: Example (left) of a hierarchy of partitions H of E, represented as (top) a sequence of partitions ordered by refinement and (bottom) its corresponding dendogram. A particular cut $\pi \in \mathit{\Pi}_E(H)$ is represented with red squared nodes.

a region $\mathcal{R}_j \in \pi_j$ such that $\mathcal{R}_i \subseteq \mathcal{R}_j$. π_i is said to refine π_j in such case. $\mathit{\Pi}_E$ is a complete lattice for the refinement (partial) ordering $\leq$. In particular, it is possible to define the refinement supremum $\pi_i \vee \pi_j$ of two partitions π_i and π_j as the smallest partition of $\mathit{\Pi}_E$ that is refined by both π_i and π_j.

A hierarchy of partitions H of E is a collection of partitions $\{\pi_i \in \mathit{\Pi}_E\}_{i=0}^n$ ordered by refinement, that is $H = \{\pi_0 \leq \pi_1 \leq \cdots \leq \pi_n\}$. π_0 is called the *leaf* partition (its regions are the *leaves* of H) and $\pi_n = \{E\}$ is the *root* of the hierarchy. A hierarchy of partitions is often represented as a tree graph (also called dendogram), where the nodes of the graph correspond to the various regions contained in the partitions of the sequence, and the vertices denote the inclusion between these regions. Alternatively, H can be equivalently defined as a collection of regions $H = \{\mathcal{R} \subseteq E\}$ that satisfy the following 3 properties:

1. $\emptyset \notin H$, $E \in H$.
2. $\forall \mathcal{R}_i, \mathcal{R}_j \in H$, $\mathcal{R}_i \cap \mathcal{R}_j \in \{\emptyset, \mathcal{R}_i, \mathcal{R}_j\}$. Any two regions belonging to H are either disjoint or nested.
3. $\forall \mathcal{R} \in H, \mathcal{R} \notin \pi_0 \Rightarrow \mathcal{R} = \bigcup_{r \in \pi_0} \{r \,|\, r \subset \mathcal{R}\}$. Any non leaf region $\mathcal{R}$ is exactly recovered by the union of all leaves of H that are included in $\mathcal{R}$.

Note that considering only items 1 and 2 allows to define tree-based representations such as the ToS, but item 3 is mandatory to define hierarchies or partitions. A *cut* of H is a partition π of E whose regions belong to H. The set of all cuts of a hierarchy H is denoted $\mathit{\Pi}_E(H)$, and is a sub-lattice of $\mathit{\Pi}_E$. All those notions related to hierarchies of partitions are summarized by Figure 1.

2.2 Braids of partitions

Braids of partitions have been introduced in [6] as a more general structure than hierarchies of partitions, and are defined as follows:

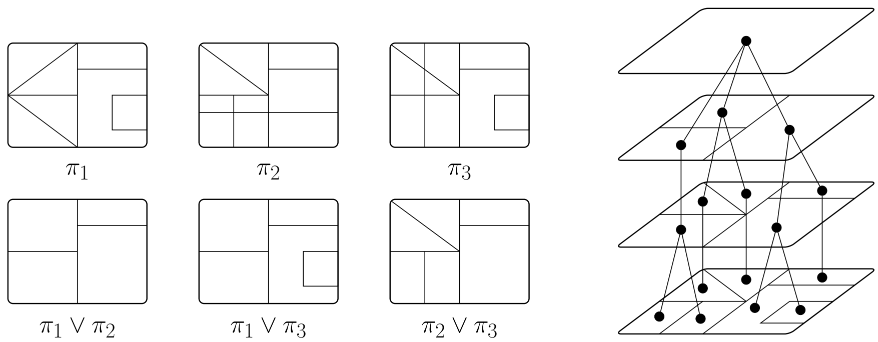

Fig. 2: Example of braid of partitions $B = \{\pi_1, \pi_2, \pi_3\}$. On the right is a monitor hierarchy of B since the pairwise refinement suprema $\pi_i \vee \pi_j, i, j \in \{1, 2, 3\}, i \neq j$ define cuts of this hierarchy different from the whole space E.

Definition 1 (Braid of partitions). *A family of partitions $B = \{\pi_i \in \Pi_E\}$ is called a braid of partitions whenever there exists some hierarchy of partitions H_m such that:*

$$\forall \pi_i, \pi_j \in B, \pi_i \vee \pi_{j \neq i} \in \Pi_E(H_m) \backslash \{E\} \tag{1}$$

Braids of partitions generalize hierarchies of partitions in the sense that the refinement ordering between the partitions composing the braid no longer needs to exist. However, there must exist some hierarchy of partition H_m, called the *monitor hierarchy*, such that the refinement supremum of any two partitions in the braid defines a cut of this hierarchy H_m. It is also worth noting that this refinement supremum must differ from the whole space $\{E\}$. Otherwise, any family of arbitrary partitions would form a braid with $\{E\}$ as a supremum, thus loosing any interesting structure. An example of braid of partitions B composed of three partitions $B = \{\pi_1, \pi_2, \pi_3\}$ as well as one possible monitor hierarchy H_m of B are displayed by Figure 2.

As we pointed out in our previous work [18], the structure of a braid B and its monitor hierarchy H_m are particulary suited for the hierarchical representation of multimodal images. As it can be observed in Figure 2, the monitor hierarchy H_m encodes all regions that are common to at least two different partitions contained in B. Assuming that the partitions composing B originate from different modalities, H_m encodes in a hierarchical manner the information that is shared by those modalities. On the other hand, all regions contained in B but not in H_m belong to a single modality, and are thus responsible for some exclusive information. Jointly considering the braid B and its monitor hierarchy H_m therefore leads to a hierarchical representation of the multimodal image that relies both on the complementary and redundant information contained in the data.

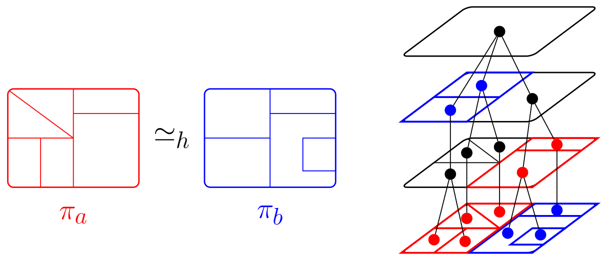

Fig. 3: Illustration of the h-equivalence relation: π_a and π_b are h-equivalent (left), they define two different cuts of the same hierarchy (right).

2.3 The h-equivalence as starting point to compose a braid

As pointed out in [6], two issues arise when working with braids of partitions:

- validating that a given family of partitions has a braid structure (that is, equation (1) is satisfied).
- defining general procedures that generate braids of partitions.

It is straightforward to compose a braid using a single hierarchy since the supremum of two cuts of a hierarchy also defines a cut of this hierarchy. For this reason, any set of cuts coming from a single hierarchy is a braid of partitions. However, this guarantee is lost when one wants to compose a braid from cuts coming from multiple hierarchies, or, even further, with arbitrary partitions (note in that respect that, although tempting to think so, the family of partitions generated by the stochastic watershed [1] has not a braid structure in general). As a matter of fact, all those cuts must be sufficiently related so their pairwise refinement suprema define cuts of the same monitor hierarchy H_m. To analyze the relationships which must be holding between the cuts of various hierarchies to form a braid, we introduced in [18] the property of *h-equivalence* (h standing for *hierarchical*):

Definition 2 (h-equivalence). *Two partitions π_a and π_b are said to be h-equivalent, and one notes $\pi_a \simeq_h \pi_b$ if and only if*

$$\forall \mathcal{R}_a \in \pi_a,\, \forall \mathcal{R}_b \in \pi_b, \mathcal{R}_a \cap \mathcal{R}_b \in \{\emptyset, \mathcal{R}_a, \mathcal{R}_b\}. \tag{2}$$

In other words, a region in π_a either refines or is a refinement of a region in π_b. Partitions π_a and π_b may not be globally comparable for the refinement ordering, but they locally are, as displayed by Figure 3. Obviously, if $\pi_a \leq \pi_b$ or $\pi_a \geq \pi_b$, then $\pi_a \simeq_h \pi_b$. All cuts of a hierarchy H are h-equivalent: $\forall \pi_1, \pi_2 \in \Pi_E(H)$, $\pi_1 \simeq_h \pi_2$. Conversely, if two partitions are h-equivalent, they define two cuts of the same hierarchy. Despite a somewhat misleading name, $\simeq_h$ is not an equivalence relation but only a tolerance relation: it is reflexive and symmetric but not transitive in general.

Following, we aim to build a braid B using cuts extracted from several hierarchical representations. To do so, we must investigate what kind of relationship must be holding between those cuts in order to guarantee the braid structure (that is, equation (1) is satisfied. Let the family of partitions $B = \{\pi_i \in \Pi_E\}$ be a braid, and let H_m be a monitor hierarchy of B.

Proposition 1. *If there exist* $\pi_i, \pi_j \in B$ *such that* $\pi_i \leq \pi_j$*, then* $\pi_j \in \Pi_E(H_m)$.

Proof. As $\pi_i \leq \pi_j$, it follows that $\pi_i \vee \pi_j = \pi_j$. And from the definition (1) of a braid, $\pi_i \vee \pi_j \in \Pi_E(H_m)$, so $\pi_j \in \Pi_E(H_m)$. □

Thus, if the braid B has two partitions ordered by refinement (two cuts extracted from the same hierarchy for instance), the coarest of them also belongs to the set of cuts $\Pi_E(H_m)$ of the monitor hierarchy H_m.

Proposition 2. *If there exist* $\pi_i, \pi_j, \pi_k, \pi_l \in B$ *such that* $\pi_i \leq \pi_j$ *and* $\pi_k \leq \pi_l$*, then* $\pi_j \simeq_h \pi_l$.

Proof. Using Proposition (1) for both $\pi_i \leq \pi_j$ and $\pi_k \leq \pi_l$, it follows that $\pi_j, \pi_l \in \Pi_E(H_m)$. Thus $\pi_j \simeq_h \pi_l$ using the property of h-equivalence. □

Therefore, if the braid B has two pairs partitions ordered by refinement, the coarest of both pairs are necessarily h-equivalent to each other since they both belong to the set of cuts $\Pi_E(H_m)$ of the monitor hierarchy H_m.

3 The braid construction procedure

Given some hierarchy H and a partition $\pi_* \in \Pi_E$, we denote by $H^{\simeq_h}(\pi_*)$ the set of cuts of H that are h-equivalent to π_*: $H^{\simeq_h}(\pi_*) \subseteq \Pi_E(H)$ with equality if and only if $\pi_* \in \Pi_E(H)$. Similarly, we denote by $H^{\leq}(\pi_*)$ the set of cuts of H that are a refinement of π_*.
Now, let H_1 and H_2 be two hierarchies of partitions built over the same space E. We aim to extract two cuts $\pi_i^1, \pi_i^2 \in \Pi_E(H_i)$ from each of those two hierarchies $H_i, i \in \{1, 2\}$ in order for the family $B = \{\pi_1^1, \pi_1^2, \pi_2^1, \pi_2^2\}$ to be a braid. For this purpose, we propose the following iterative procedure:

1. First select arbitrarily some cut $\pi_1^1 \in \Pi_E(H_1)$.
2. Then choose a cut π_2^1 in the constrained set $H_2{}^{\simeq_h}(\pi_1^1) \backslash \{E\}$, that is, a cut from H_2 which is h-equivalent to π_1^1 and different from the whole space $\{E\}$.
3. Finally, complete by taking a cut in each hierarchy that is a refinement of the cut previously extracted from the other hierarchy, that is $\pi_i^2 \in \Pi_E(H_i), i \in \{1, 2\}$ such that $\pi_1^2 \leq \pi_2^1$ and $\pi_2^2 \leq \pi_1^1$.

This procedure is summarized by Figure 4.

Proposition 3. *Under this configuration,* $B = \{\pi_1^1, \pi_1^2, \pi_2^1, \pi_2^2\}$ *has a braid structure.*

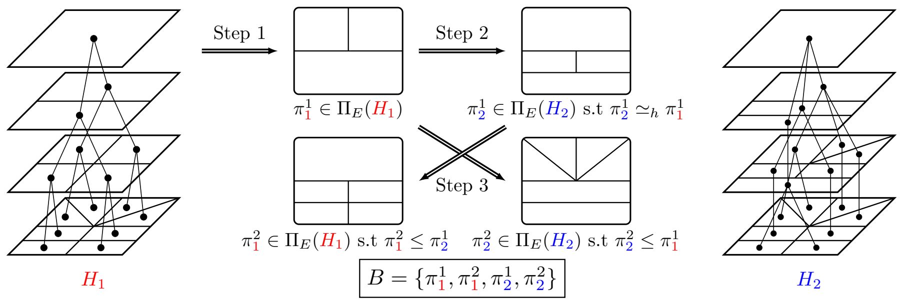

Fig. 4: Composing a braid B with cuts from two hierarchies H_1 and H_2.

Proof. Let $B = \{\pi_1^1, \pi_1^2, \pi_2^1, \pi_2^2\}$ be a family of partitions composed following the previously described procedure, and let $\pi_{i,j}^{k,l} = \pi_i^k \vee \pi_j^l$ denote the pairwise refinement suprema of partitions in B. In particular, the 4 partitions composing B generates $\binom{4}{2} = 6$ different pairwise refinement suprema $\pi_{1,1}^{1,2}, \pi_{1,2}^{1,1}, \pi_{1,2}^{1,2}, \pi_{1,2}^{2,1}, \pi_{1,2}^{2,2}, \pi_{2,2}^{1,2}$. Checking that B is a braid amounts to verify whether the $\pi_{i,j}^{k,l}$ all defines cuts of the same monitor hierarchy H_m, which is equivalent to showing that they are (at least) all h-equivalent to each other. In order to show the braid structure of B, we first demonstrate the following result:

Lemma 1. *Let $\pi_1, \pi_2, \pi_3 \in \Pi_E$ be some partitions of E such that $\pi_1 \simeq_h \pi_3$ and $\pi_2 \leq \pi_3$. Then $\pi_1 \vee \pi_2 \simeq_h \pi_3$.*

Proof. If $\pi_1 \leq \pi_3$, then $\pi_1 \vee \pi_2 \leq \pi_3$ by definition of the refinement supremum, and so $\pi_1 \vee \pi_2 \simeq_h \pi_3$ since any two ordered partitions are also h-equivalent.
On the other hand, if $\pi_1 \geq \pi_3$, then $\pi_1 \geq \pi_2$, hence $\pi_1 \vee \pi_2 = \pi_1$ and so $\pi_1 \vee \pi_2 \simeq_h \pi_3$ for the same reason as above.
In the most general case where π_1 and π_3 are h-equivalent but can nonetheless not be ordered, it means that π_1 is a refinement of π_3 in some parts of E, and is refined by π_3 in the other parts. In the former case, let $\mathcal{R}_3$ be a region of π_3 and $\pi_1(\mathcal{R}_3), \pi_2(\mathcal{R}_3)$ be the refinements (partial partitions) of $\mathcal{R}_3$ in π_1 and π_2. Then, $\pi_1(\mathcal{R}_3) \vee \pi_2(\mathcal{R}_3)$ is also a refinement of $\mathcal{R}_3$, implying that $\pi_1 \vee \pi_2$ refines π_3 in the part of E covered by $\mathcal{R}_3$. In the case where π_3 is locally a refinement of π_1, then given $\mathcal{R}_1 \in \pi_1$, there exists a refinement $\pi_3(\mathcal{R}_1)$ of $\mathcal{R}_1$ in π_3, and therefore a refinement $\pi_2(\mathcal{R}_1)$ of $\mathcal{R}_1$ in π_2 since $\pi_2 \leq \pi_3$. Therefore, $\{\mathcal{R}_1\} \vee \pi_2(\mathcal{R}_1) = \{\mathcal{R}_1\}$ and thus π_3 refines $\pi_1 \vee \pi_2$ in the part of E covered by $\mathcal{R}_1$. Finally, $\pi_1 \vee \pi_2$ either refines or is refined by π_3 in all parts of E, hence $\pi_1 \vee \pi_2 \simeq_h \pi_3$. □

To ease the reading of the proof of Proposition (3), we first recall the relations holding between the various partitions composing the braid B:

- $\pi_1^1 \simeq_h \pi_2^1$ by construction.
- $\pi_1^2 \leq \pi_2^1$ and $\pi_2^2 \leq \pi_1^1$ by construction.
- $\pi_1^1 \simeq_h \pi_1^2$ beacuse they are both cuts of the same hierarchy H_1. Similarly, $\pi_2^1 \simeq_h \pi_2^2$.

Following, we prove that all the pairwise refinement suprema of B are at least all h-equivalent to each other. Their relationships are summarized in table 1.

Table 1: Summary of the relationships holding between all pairwise refinement suprema of B with their corresponding item in the proof.

	$\pi^{1,2}_{1,1}$	$\pi^{1,1}_{1,2}$	$\pi^{1,2}_{1,2}$	$\pi^{2,1}_{1,2}$	$\pi^{2,2}_{1,2}$	$\pi^{1,2}_{2,2}$
$\pi^{1,2}_{1,1}$	X	1. $\leq$	2. $\leq$	3. $\simeq_h$	4. $\leq$	5. $\simeq_h$
$\pi^{1,1}_{1,2}$		X	6. $\leq$	7. $\leq$	8. $\leq$	9. $\leq$
$\pi^{1,2}_{1,2}$			X	10. $\simeq_h$	11. $\simeq_h$	12. $\simeq_h$
$\pi^{2,1}_{1,2}$				X	13. $\simeq_h$	14. $\leq$
$\pi^{2,2}_{1,2}$					X	15. $\leq$
$\pi^{1,2}_{2,2}$						X

1. $\pi^{1,2}_{1,1} = \pi^1_1 \vee \pi^2_1$. As $\pi^2_1 \leq \pi^1_2$ by construction of B, it follows that $\pi^1_1 \vee \pi^2_1 \leq \pi^1_1 \vee \pi^1_2$, hence $\pi^{1,2}_{1,1} \leq \pi^{1,1}_{1,2}$.
2. $\pi^{1,2}_{1,2} = \pi^1_1 \vee \pi^2_2 = \pi^1_1$ as $\pi^2_2 \leq \pi^1_1$ by construction of B. By property of the refinement supremum, one has $\pi^1_1 \leq \pi^1_1 \vee \pi^2_1 = \pi^{1,2}_{1,1}$, hence $\pi^{1,2}_{1,2} \leq \pi^{1,2}_{1,1}$.
3. By construction of B, one has $\pi^1_1 \simeq_h \pi^1_2$ and $\pi^2_1 \leq \pi^1_2 = \pi^{2,1}_{1,2}$. Using lemma 1, it follows that $\pi^1_1 \vee \pi^2_1 = \pi^{1,2}_{1,1} \simeq_h \pi^{2,1}_{1,2}$.
4. $\pi^2_2 \leq \pi^1_1$ by construction of B, meaning that $\pi^2_1 \vee \pi^2_2 \leq \pi^2_1 \vee \pi^1_1$, hence $\pi^{2,2}_{1,2} \leq \pi^{1,2}_{1,1}$.
5. Using item 3, we first have $\pi^{1,2}_{1,1} \simeq_h \pi^1_2 = \pi^{2,1}_{1,2}$. In addition, $\pi^2_2 \leq \pi^1_1$ by construction of B, implying that $\pi^2_2 \leq \pi^1_1 \vee \pi^2_1 = \pi^{1,2}_{1,1}$. Using lemma 1 finally leads to $\pi^{1,2}_{1,1} \simeq_h \pi^{1,2}_{2,2}$.
6. $\pi^{1,2}_{1,2} = \pi^1_1$ as $\pi^2_2 \leq \pi^1_1$ by construction of B. The basic property of the refinement supremum allows to conclude that $\pi^1_1 \leq \pi^1_1 \vee \pi^1_2$, hence $\pi^{1,2}_{1,2} \leq \pi^{1,1}_{1,2}$.
7. The exact same reasoning as item 6 applied to $\pi^{2,1}_{1,2} = \pi^1_2$ leads to $\pi^{2,1}_{1,2} \leq \pi^{1,1}_{1,2}$.
8. $\pi^2_1 \leq \pi^1_2$ and $\pi^2_2 \leq \pi^1_1$, both by construction of B. It immediately follows that $\pi^2_1 \vee \pi^2_2 \leq \pi^1_1 \vee \pi^1_2$, hence $\pi^{2,2}_{1,2} \leq \pi^{1,1}_{1,2}$.
9. The same reasoning as item 1 applies to $\pi^{1,2}_{2,2} = \pi^1_2 \vee \pi^2_2$, leading to $\pi^{1,2}_{2,2} \leq \pi^{1,1}_{1,2}$.
10. By construction of B, one has $\pi^1_1 = \pi^{1,2}_{1,2} \simeq_h \pi^{2,1}_{1,2} = \pi^1_2$, hence the result.
11. $\pi^{1,2}_{1,2} = \pi^1_1$ as $\pi^2_2 \leq \pi^1_1$ by construction of B. In addition, $\pi^1_1 \simeq_h \pi^2_1$ as they are both cuts of the same hierarchy H_1. Using lemma 1, it follows that $\pi^1_1 = \pi^{1,2}_{1,2} \simeq_h \pi^{2,2}_{1,2} = \pi^2_1 \vee \pi^2_2$.
12. The same reasoning as item 3 applies to $\pi^{1,2}_{2,2}$ and $\pi^{1,2}_{1,2} = \pi^1_1$ and, relying upon lemma 1, leads to $\pi^{1,2}_{1,2} \simeq_h \pi^{1,2}_{2,2}$.

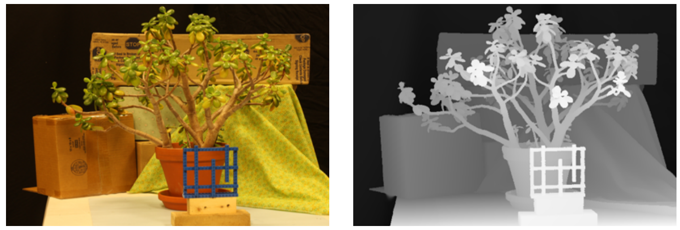

(a) (b)

Fig. 5: (a) RGB modality and (b) Depth map of a RGB/Depth multimodal image from the Middlebury Stereo Dataset.

13. The same reasoning as item 11 applies to $\pi^{2,1}_{1,2} = \pi^1_2$ and $\pi^{2,2}_{1,2}$, leading to $\pi^{2,1}_{1,2} \simeq_h \pi^{2,2}_{1,2}$.
14. The same reasoning as item 2 applies to $\pi^{1,2}_{2,2}$ and $\pi^{2,1}_{1,2} = \pi^1_2$, leading to $\pi^{2,1}_{1,2} \leq \pi^{1,2}_{2,2}$.
15. The same reasoning as item 4 applies to $\pi^{2,2}_{1,2}$ and $\pi^{1,2}_{2,2}$, leading to $\pi^{2,2}_{1,2} \leq \pi^{1,2}_{2,2}$.

Finally, all the pairwise refinement supremum $\pi^{k,l}_{i,j} = \pi^k_i \vee \pi^l_j$ that can be formed using the partitions belonging to B are (at least) all h-equivalent to each other. Therefore, there exists some hierarchy H_m such that all $\pi^{k,l}_{i,j} \in \Pi_E(H_m)$, which proves that B has a braid structure when constructed following the proposed procedure. □

4 A quick look into the braid structure

The braid construction procedure presented in Section 3 and summarized by Figure 4 is only operable when two cuts are extracted from two hierarchies of partitions. While this may appear quite restrictive from an applicative point of view, we are up to now only able to provide the proposed procedure in this context. As a matter of fact, the braid structure is guaranteed whenever the refinement suprema of all partitions composing the braid are all h-equivalent to each other, but some additional efforts are still needed to understand how this h-equivalence constraint can be formulated as a constraint on the partitions of the braid.

In order to give a brief insight on the structural properties of the braid of partitions and its monitor hierarchy when the proposed procedure is applied on a real multimodal scenario, we consider the RGB/Depth multimodal image presented by Figure 5, originating from the 2014 database of the Middlebury Stereo Dataset [15]. Both modalities are co-registred and comprise 400×300 pixels.

The two hierarchies of partitions H_1 and H_2 are obtained by means of two BPTs [13] with standard parameters (mean value for the region model and Euclidean distance for the merging criterion). We use the same initial partition

Table 2: Number of regions $NbReg$ in all partitions composing the braid B when the first cut is extracted from H_1 built on RGB modality.

π_1^1	10	20	50	100	200	500	1000
$\pi_2^1 \simeq_h \pi_1^1$	544	822	1095	1267	1415	1609	1684
$\pi_1^2 \leq \pi_2^1$	1358	1527	1680	1750	1800	1860	1876
$\pi_2^2 \leq \pi_1^1$	544	822	1095	1267	1415	1620	1801

Table 3: Number of regions $NbReg$ in all partitions composing the braid B when the first cut is extracted from H_1 built on Depth modality.

π_1^1	10	20	50	100	200	500	1000
$\pi_2^1 \simeq_h \pi_1^1$	633	829	994	1138	1315	1569	1699
$\pi_1^2 \leq \pi_2^1$	1530	1612	1657	1707	1785	1864	1865
$\pi_2^2 \leq \pi_1^1$	633	829	994	1138	1315	1578	1824

for both BPTs, namely the intersection of two mean shift clustering procedures ran independently on each modality (yielding a total of 2148 leaf regions). The first cut π_1^1 is extracted from H_1 as the partition composed of the $NbReg$ last regions before completion of the region merging process of the BPT. It defines the extraction of the three remaining cuts $\pi_1^2, \pi_2^1, \pi_2^2$ through the application of the procedure presented in Section 3 as follows: $\pi_2^1 = \bigvee\{H_2{}^{\simeq_h}(\pi_1^1)\backslash\{E\}\}$, $\pi_1^2 = \bigvee\{H_1{}^{\leq}(\pi_2^1)\}$ and $\pi_2^2 = \bigvee\{H_2{}^{\leq}(\pi_1^1)\}$. The three constrained set of cuts $H_2{}^{\simeq_h}(\pi_1^1)\backslash\{E\}$, $H_1{}^{\leq}(\pi_2^1)$ and $H_2{}^{\leq}(\pi_1^1)$ are non empty since both hierarchies H_1 and H_2 have the same leaf partition (all three sets thus contain at least the leaf partition). We set $NbReg$ (that is, the number of regions in π_1^1) to $10, 20, 50, 100, 200, 500$ and 1000. Tables 2 and 3 give the number of regions composing partitions π_1^2, π_2^1 and π_2^2 when H_1 is the BPT built on the RGB and Depth modality, respectively. As it can be seen, there is in both cases a big gap between the number of regions in π_1^1 and π_2^1. Even though π_2^1 is defined to be h-equivalent to π_1^1, it turns out in practice to be a refinement of π_1^1 since $\pi_2^2 \leq \pi_1^1$ has the same number of regions as π_2^1 (except for $NbReg = 500, 1000$). The reason is that when π_1^1 has a relatively small number of regions, there are no regions in H_2 that are refined by those of π_1^1, hence the largest cut in $H_2{}^{\simeq_h}(\pi_1^1)\backslash\{E\}$ is equivalent to the one in $H_2{}^{\leq}(\pi_1^1)$. There is an equivalently big gap between the number of regions in π_2^1 and π_1^2. While this can be explained by the fact that π_1^2 is defined as a refinement of π_2^1, it also means that it is very difficult to find "intermediary" regions in H_1 (that is, regions that are not close from the root or the leaves of H_1, and that are more likely to be associated with semantic objects in the image) that correspond to those in H_2. While this is out of the scope of this paper, this interpretation might be a potential issue in a practical scenario of braid-based hierarchical representation of multimodal images and will be further investigated.

Regarding the influence of the hierarchy from which is extracted the first cut

π_1^1, it can be appreciated by comparing Table 2 and Table 3 that the proposed procedure yields cuts π_1^2, π_2^1 and π_2^2 whose number of regions remains relatively stable whether H_1 was defined to be the BPT constructed on the RGB modality or the Depth modality. This observation will have to be validated with more in-depth experiments, and the influence of the choice of the modality associated with H_1 on practical image processing tasks (such as segmentation or object recognition) will also need to be evaluated (since the semantic content of those regions will depend on the identity of the modality associated with H_1 and H_2). It could also be interesting to study whether the number of regions $NbReg$ of the first partition π_1^1 can vary within a given range of values without impacting the structure of the monitor hierarchy. Similarly, the influence of the choice of π_1^2, π_2^1 and π_2^2 in their respective constrained sets of cuts will also have to be investigated. While these considerations are probably data dependent, they could nevertheless give a deeper insight on the stability of the braid structure generated by the proposed procedure.

5 Conclusion

Braids of partitions were defined as a generalization of hierarchies of partitions, in the sense that all partitions composing the braid do not need to be ordered by refinement. This more permissive property opens the door to potentially several applications of interest for the braid structure, but it also brings difficulties to build such structure in practice. In our previous work [18], we experimentally provided a procedure to build a braid of partitions as a combination of cuts coming from hierachies, and intuited the potential of braids to perform multimodal image segmentation.

In this paper, we formally demonstrated that the procedure proposed in [18] was indeed yielding a braid of partitions. While we did not focus here on a more thorough evaluation of the usefulness of the braid structure for multimodal image analysis, this is obviously an important future research avenue. In addition, we are up to now bound to build a braid by using only two hierarchies of partitions, and two cuts per hierarchy. In that respect, future work will investigate theoretical aspects related to the construction of the braid of partitions, namely how to extract more cuts coming from various hierarchies and still maintain the braid structure.

References

1. J. Angulo and D. Jeulin. Stochastic watershed segmentation. In *PROC. of the 8th International Symposium on Mathematical Morphology*, pages 265–276, 2007.
2. S. Beucher. Watershed, hierarchical segmentation and waterfall algorithm. In *Mathematical morphology and its applications to image processing*, pages 69–76. Springer, 1994.
3. P. Bosilj, E. Kijak, and S. Lefèvre. Partition and inclusion hierarchies of images: A comprehensive survey. *Journal of Imaging*, 4(2):33, 2018.

4. E. Carlinet and T. Géraud. MToS: A tree of shapes for multivariate images. *IEEE Transactions on Image Processing*, 24(12):5330–5342, 2015.
5. J. Cousty, L. Najman, Y. Kenmochi, and S. Guimarães. Hierarchical segmentations with graphs: quasi-flat zones, minimum spanning trees, and saliency maps. *Journal of Mathematical Imaging and Vision*, 60(4):479–502, 2018.
6. B. R. Kiran and J. Serra. Braids of partitions. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 217–228. Springer, 2015.
7. P. Monasse and F. Guichard. Fast computation of a contrast-invariant image representation. *Image Processing, IEEE Transactions on*, 9(5):860–872, 2000.
8. L. Najman and J. Cousty. A graph-based mathematical morphology reader. *Pattern Recognition Letters*, 47:3–17, 2014.
9. G. Palou and P. Salembier. Hierarchical video representation with trajectory binary partition tree. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2099–2106, 2013.
10. N. Passat and B. Naegel. Component-trees and multivalued images: Structural properties. *Journal of Mathematical Imaging and Vision*, 49(1):37–50, 2014.
11. B. Perret, J. Cousty, S. J. F. G. aes, and D. S. Maia. Evaluation of hierarchical watersheds. *IEEE Transactions on Image Processing*, 27(4):1676–1688, 2018.
12. J. F. Randrianasoa, C. Kurtz, E. Desjardin, and N. Passat. Binary partition tree construction from multiple features for image segmentation. *Pattern Recognition*, 84:237–250, 2018.
13. P. Salembier and L. Garrido. Binary partition tree as an efficient representation for image processing, segmentation, and information retrieval. *Image Processing, IEEE Transactions on*, 9(4):561–576, 2000.
14. P. Salembier, A. Oliveras, and L. Garrido. Antiextensive connected operators for image and sequence processing. *Image Processing, IEEE Transactions on*, 7(4):555–570, 1998.
15. D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German Conference on Pattern Recognition*, pages 31–42. Springer, 2014.
16. P. Soille. Constrained connectivity for hierarchical image partitioning and simplification. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(7):1132–1145, 2008.
17. G. Tochon. *Hierarchical analysis of multimodal images*. PhD thesis, Université Grenoble Alpes, 2015.
18. G. Tochon, M. Dalla Mura, and J. Chanussot. Segmentation of multimodal images based on hierarchies of partitions. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 241–252. Springer, 2015.
19. V. Vilaplana, F. Marques, and P. Salembier. Binary partition trees for object detection. *IEEE Transactions on Image Processing*, 17(11):2201–2216, 2008.
20. Y. Xu, T. Géraud, and L. Najman. Connected filtering on tree-based shape-spaces. *IEEE transactions on pattern analysis and machine intelligence*, 38(6):1126–1140, 2016.
21. Y. Xu, T. Géraud, and L. Najman. Hierarchical image simplification and segmentation based on mumford–shah-salient level line selection. *Pattern Recognition Letters*, 83:278–286, 2016.