Subsymbolic, hybrid and explainable AI
What can it change in medical imaging?

Isabelle Bloch

1. Introduction

While symbolic methods and statistical machine learning methods for artificial intelligence (AI) have been developing rather independently for decades, with alternated predominance of one or the other across time, a current trend is to merge both types of approaches. Examples include neuro-symbolic approaches (see e.g., De Raedt et al. 2020; d'Avila Garcez/Lamb 2023; Garnelo/Shanahan 2019; Kautz 2022; Marcus 2020), among others. However, in this paper, hybrid artificial intelligence is intended in a broader sense, as the combination of several AI methods, whatever their type.[1] These methods may belong to the domains of abstract knowledge representation and formal reasoning, based on logic, structural representation (such as graphs and hypergraphs, ontologies, concept lattices, etc.), machine learning, etc. Additionally, imprecision in data, knowledge and reasoning can benefit from the fuzzy sets theory.

Such combinations of approaches take inspiration from cognitive functions. Roughly speaking, according to Kahneman (2012), who distinguished two systems for thinking named system 1 and system 2, we may consider, from a (strongly simplified) AI point of view, modeling system 1 (rapid, intuitive) by deep learning and system 2 (slower, more controlled, logical) by symbolic reasoning. Developing neuro-symbolic approaches is a new trend to combine the two systems (see e.g., Kautz 2022). But again, more theories will be committed in our view of hybrid AI, in particular for image understanding.

The aim of this paper is not to propose new methods for hybrid AI, but rather, as a position paper, to highlight how this way of thinking and design-

1 We should note here that AI is already the umbrella term for very different methods, and that many AI methods or systems are actually by essence hybrid.

ing AI systems offers opportunities towards explainability in the field of explainable AI (XAI) and as a mean to maintain the link between knowledge and data. In that domain, too, the two main branches are developed quite independently, with early work (e.g., Peirce at the end of the 19th century) focusing on logical reasoning based on abduction on the one hand, versus recent methods focusing on features or data most involved in a decision on the other hand (to name but a few). In the first paradigm, knowledge is represented by symbols in a given logic and the reasoning power of this logic then plays a major role. Reasoning is based on axioms, theories and inference rules, leading to provable, non-refutable conclusions. In the second paradigm, where data and experience play the major role, statistical guarantees can be achieved, but conclusions are potentially refutable. As an example, fuzzy sets can cope with both approaches and establish links between them.

These ideas are illustrated in the field of image understanding and formulated as a spatial reasoning problem (section 2). Examples of combinations of different AI methods are given, both for knowledge and data representation, in section 3, and for reasoning in section 4. These methods find concrete applications in several domains such as medical imaging (only briefly mentioned in this paper). The question of explanations is addressed in section 5. Finally a short discussion on open research directions concludes the paper (section 6).

This paper is an extension of Bloch (2022), and focuses on the explainability aspects as well as the usefulness of hybrid AI and XAI for medical image understanding, in particular in pediatrics. The example of pediatric imaging is relevant here for illustrating the main topics developed in this paper, because of the challenging issues it raises (few data, very specific images, anatomy and pathologies, etc.). In addition, as mentioned in the next section, it is important with regards to the availability of domain knowledge and the usefulness of developing tools for explainable image understanding. This paper does not contain technical details – those can be found in the listed references.

2. Image understanding and spatial reasoning

Image understanding, at the simplest level, refers to the problem of recognizing an object or structure, or several objects in an image, which can either be real, as an observation of a part of the real world, or synthetic. But this may not be sufficient and more generally, relations between these objects should be considered towards a global recognition of the scene and a higher level interpretation, beyond individual objects. Furthermore, the recognition of an individual object can benefit from the recognition of others.

The question of semantics is central, since it is not directly in the image, but should be inferred based on visual features. We advocate that knowledge should be involved in this process. Indeed, while purely data driven approaches have proven powerful in image and computer vision problems, with sometimes impressive results, they still require a good accessibility to numerous and annotated data, where annotations bring the semantic information. This is not always possible and induces high costs (in terms of both human interactions and computation). Knowledge and models have then an important role to play. Image understanding is formulated as a spatial reasoning problem, combining representations of data and knowledge, pertaining to both objects and relations between objects (in particular spatial relations), as well as reasoning on them.

Let us take the example of pediatric medical imaging. In this domain, data may be scarce and present a high variability. Data are also very heterogeneous when they come from multicentric studies, with different hospitals, different imaging machines, different protocols and acquisition parameters. This makes the appearance of the same tissues, organs or pathologies vary a lot from one image to the other. This problem is sometimes addressed by transferring a model learned on adult images to children images. However, there is a huge domain gap, since the relative sizes of body parts, organs and pathologies vary considerably (in particular depending on the development stage of the children). Pathologies of children may differ from those observed in adults, the acquisitions should be as short as possible on children, thus inducing differences in image appearance. The contrast between tissues can also be quite different, even with the same acquisition protocol. Control cases and images of healthy children are even more rare, in particular due to ethical reasons. All this makes the problem particularly difficult. On the other hand, anatomical and medical knowledge is important, and was gathered over centuries. Using it is undoubtedly helpful.

Spatial reasoning has been largely developed in symbolic AI, based mostly on logic and benefitting from the reasoning apparatus of this logic (Aiello/Pratt-Hartmann/Benthem 2007). It has been much less developed for image understanding, where purely symbolic approaches are limited to account for numerical information. This again votes for hybrid approaches. Spatial reasoning evolved from purely qualitative and symbolic approaches, to more and more hybrid methods involving methods from mathematical morphology, fuzzy sets, graphs, machine learning, etc. to gain in expressivity (sometimes at the price of increased complexity). As an example, let us mention region connection calculus (RCC), that was first proposed in logical frameworks (first order, modal) and then augmented with fuzzy sets to handle imprecision, with mathematical morphology, lattice-based reasoning, etc. (Aiello/Pratt-Hartmann/Benthem 2007; Aiguier/Bloch 2019; Bloch 2021b; Landini et al. 2019; Randell/Cui/Cohn 1992; Schockaert et al. 2008; Schockaert/De Cock/Kerre 2009). The main ingredients in spatial reasoning include knowledge representation, imprecision representation and management, fusion of heterogeneous information (whether it is knowledge or data), reasoning and decision making. Approaches for spatial reasoning take a lot of inspiration from work in philosophy, linguistics, human perception, cognition, neuroimaging, art, etc. (see e.g., a related discussion for the case of spatial distances in Bloch 2003).

Models for image understanding are particularly useful to represent, in a formal way, knowledge (about the domain, the scene content and in particular its structure), image information (type of acquisition, geometry, characteristics of signal and noise, etc.), the potential imperfections of knowledge and data (imprecision, uncertainty, incompleteness, etc.), as well as the combination of knowledge and image information. These models are then included in algorithms to guide image understanding in concrete applications. Conversely, models can be built from data, to infer knowledge, or to provide a digital twin of a patient as a 3D model, useful to plan a surgery or a therapy, as well as to explain the plan (e.g., to other surgeons, to the patients and their parents in the case of pediatrics).

An important issue is the semantic gap (Smeulders et al. 2000), with the following question: how to link visual percepts from the images to symbolic descriptions? In artificial intelligence, this is close to the notions known as the anchoring or symbol grounding problem (Coradeschi/Saffiotti 1999; Harnad 1990). Solving the semantic gap issue has bidirectional consequences: on the one hand, it allows moving from a concept to its instantiation in the image (or feature) space, as a guide during spatial reasoning. On the other hand, it is part of the explainability, since it links results inferred from the image to concepts related to prior knowledge. For instance, anatomical knowledge says that the heart is between the lungs. Since the heart might be difficult to recognize directly in a medical image (e.g., a non-enhanced CT image), we may rely on its relative position with respect to the lungs (which are easier to detect in such images) to perform the task. This is an example where the recognition of an object benefits from the recognition of other objects, as mentioned at the beginning of this section. Conversely, we can explain the recognition of an image region as the heart because it is between the lungs (see section 5).

3. Information and knowledge representation

Representations of spatial entities can take various forms, either in the spatial domain (region, key points, bounding box, etc.), or abstractly, as in region connection calculus (RCC), as formulas in a given logic. Semi-quantitative (or semi-qualitative) representations as fuzzy sets (in either domain) constitute a good midway and can accommodate both numerical and symbolic representations (Zadeh 1965). Representations as numbers, imprecise numbers, intervals, distributions and linguistic values can all find a unifying framework with fuzzy sets. In this framework, different types of imperfections can be easily modeled, such as imprecision or blurriness on the boundaries of an object, on its location, shape or appearance, ambiguity, partial lack of information, etc. These imperfections can have varied sources, starting with the observed phenomenon, the sensors and the associated image reconstruction algorithms, and can also result from image processing steps such as filtering, registration and segmentation.

Spatial reasoning involves models of spatial entities, but also spatial relations between these entities. Here, the advantages of fuzzy representations become even more significant. This was already stated in the 1970s (Freeman 1975), but formal mathematical models were developed only later (see the review in Bloch 2005). The objective is to account for the intrinsic imprecision of concepts such as "close to", "to the left of" and "between", which are nevertheless perfectly understandable by humans in a given context and to account for the imprecision of the objects (even for a conceptually well-defined relation). In our previous work, we have designed mathematical models of several relations (set theoretical, topological, distances, directional relations and more complex relations such as between, along, parallel, etc.) by combining formalisms from mathematical morphology and fuzzy sets. They are detailed in Bloch and Ralescu (2023), chapter 6, and in the references cited therein. From a mathematical point of view, the common underlying structure is the one of complete lattices that allows instantiating the definitions, with the very same formalism in different frameworks: sets, fuzzy sets, graphs and hypergraphs, formal concept lattices, conceptual graphs, ontologies, etc., that can all be endowed with a lattice structure with appropriate partial orders. This becomes particularly useful when defining spatial relations based on mathematical morphology, a theory where deterministic operators are usually defined in a lattice. Our main idea was to design structuring elements, defined as fuzzy sets in the spatial domain, that provide the semantics of the spatial relation. Then applying a fuzzy morphological dilation of a reference object (whether fuzzy or not) using this structuring element provides the region of space where the considered relation is satisfied. The membership value of a point to the resulting fuzzy set is then interpreted as the degree to which the relation of this point to the reference object is satisfied. This approach can be applied to several classes of spatial relations: topological, distances, relative direction and more complex ones such as along, parallel, between, etc. (see e.g., Bloch 2021a; Bloch/Ralescu 2023 and the references therein). It applies to objects defined as sets or fuzzy sets in the spatial domain, but also those defined more abstractly as logical formulas, vertices of a (hyper-)graph, concepts, etc.

Note that most of the frameworks mentioned above carry structural information, useful for instance when representing the spatial arrangement of objects in a scene and in an image. To take a simple example, a graph can represent this structure, where vertices correspond to objects (e.g., anatomical structures in medical images) and edges correspond to relations between objects (e.g., contrast between two structures in a given imaging modality, relative position between objects, etc.), this graph being enhanced with the fuzzy representations of objects and their properties, as well as relations. For instance, the representation of a spatial relation can be abstract, as extracted from an ontology for example, or linked to the concrete domain of an image (degree of satisfaction of the relation, region of space where the relation to some object is satisfied, etc.), using linguistic variables, as explained next. Other structured representations of knowledge (including spatial knowledge) may rely on grammars, decision trees, relational algebras, or on temporal or spatial configurations and graphical models. They can also benefit from a fuzzy modeling layer, helping them cope with imprecision.

The relevance of fuzzy sets for knowledge representation, combined with other representations, lies in their ability to capture linguistic as well as quantitative knowledge and information. A useful notion is the one of linguistic variable (Zadeh 1975), where symbolic values, defined at an ontological level, have semantics defined by membership functions on a concrete domain at the image or features level. The membership functions and their parameters can be handcrafted, according to some expert knowledge on the application domain. They can also be learned, for instance from annotated data (Atif et al. 2007). The advantage of such representations is that linguistic characterizations may be less specific than numerical ones (and therefore need less information). Their two levels (syntactic and semantic) allow on the one hand for approximate modeling of vague concepts, and reasoning on them, and on the other hand constitute an efficient way to solve the semantic gap issue (see section 2) by providing semantics in concrete domains, according to each specific context. Linguistic variables, maintaining the consistency between concepts and data, therefore play an important role for explainability. Similarly, the goals of an image understanding problem can be expressed in an imprecise way, and again, translating vague concepts into useful representations and algorithms benefits from fuzzy modeling, in particular when using linguistic variables.

4. Reasoning

Based on the previous representations, the reasoning part takes various forms, separately or in combination, again in the spirit of hybrid AI. It is important to mention a few, mostly from previous work, which led to applications in medical imaging, in particular for brain structure recognition:[2] matching between a model and an image based on graph representations (Aldea/Bloch 2010; Cesar et al. 2005; Fasquel/Delanoue 2019; Perchant/Bloch 2002); sequential spatial reasoning mimicking the usual cognitive process where one may focus on an object that is easy to detect and to recognize, and then move progressively to more and more difficult objects by exploring the space based on the spatial relations with respect to previously recognized objects (Bloch/Géraud/Maître 2003; Colliot/Camara/Bloch 2006; Delmonte et al. 2019; Fouquier/Atif/Bloch 2012); exploration of the whole space and reducing progressively the potential region for each object, again mimicking a type of cognitive process, for instance by expressing the task as a constraint satisfaction problem (Deruyver/Hodé 1997; Nempont/Atif/Bloch 2013), logical reasoning based on abduction, to find the best explanations to the observations according to the available knowledge (Yang/Atif/Bloch 2015) and logical reasoning driven by an ontology (Hudelot/Atif/Bloch 2008).

2 These are only examples and similar approaches have been developed in other application domains, such as satellite imaging, video, music representations, etc.

In all these methods, an important feature is the combination of several approaches within the framework of hybrid AI, with the aim of explainability. Abstract knowledge representation and formal reasoning (typically using logics) are appropriate to build a knowledge base representing prior information (on anatomy for the considered examples) and to reason on it – the expressivity and the reasoning power depending on the chosen logic. Structural representations (graphs and hypergraphs, ontologies, conceptual graphs, concept lattices, etc.) are frameworks to convert expert knowledge on the spatial organization of objects (e.g., organs in medical imaging) into operational computational models. As mentioned in section 3, converting knowledge into meaningful representations and algorithms highly benefits from fuzzy modeling, in particular linguistic variables used to fill the semantic gap. This is indeed key to explainability. These models are then associated with structural representations to enrich them. For instance, fuzzy models of object features (shape, appearance) and of spatial relations can be attributes of vertices or edges of graphs, associated with concept descriptions in ontologies or conceptual graphs, providing semantics for these concepts, and considering them properties in fuzzy extensions of concept lattices, or providing semantics of logical formulas.

Usually several pieces of knowledge are involved together in the reasoning process. The advantages of fuzzy sets lie in the variety of combination operators, offering a lot of flexibility in their choice, that can be adapted to any situation at hand, and which may deal with heterogeneous information (Dubois/Prade 1985; Yager 1991). A classification of these operators was proposed by Bloch (1996), with respect to their behavior (in terms of conjunctive, disjunctive, compromise (Dubois/Prade 1985), the possible control of this behavior, their properties and their decisiveness.

Now, considering the recent huge developments in machine learning, and in particular deep learning, a recent trend is to combine such approaches with knowledge driven methods. This can be done at several levels (see e.g., Xie et al. 2021): to enhance the input (e.g., by including in the input of a neural network as a result of some image processing method as in Couteaux et al. 2019), as regularization terms in the loss function (e.g., to force the satisfaction of some relations), or to focus attention on specific patches based on geometric or topological information (e.g., vessel tree, see Virzi et al. 2018), or as post-processing to improve results (e.g., Chopin et al. 2022). Conversely, in some situations, the neural networks can use implicit spatial relations to solve a task such as object segmentation and recognition, as soon as the concerned objects are within the receptive field (Riva et al. 2022). Again, one of the advantages of such hybrid approaches is to improve interpretability and explainability. This is particularly important in medical imaging for increasing the confidence the user may have in an approach based on deep learning, consequently also increasing the adoption of such techniques.

Finally, the result of an image understanding system can be expressed in various forms (sets of (fuzzy) objects representing recognized structures, classes (of objects or pathologies for instance), properties of objects or structures and the relations they share, linguistic descriptions providing in a given vocabulary sentences describing the content of the image, etc.), finding yet again a unifying representation framework in fuzzy sets. The next step is then to provide explanations to these results.

5. Explanations

A first way to provide explanations is to rely on abductive reasoning in some logic.[3] Mathematical morphology is a useful theory for abductive reasoning and various logics (Aiguier et al. 2018; Bloch 2006; Bloch et al. 2018). An example is the use of erosion or derived operators to provide explanations for observations according to a knowledge base by applying these operators to a set of models for logical formulas or to a concept lattice. For instance, from a knowledge base on anatomy, expressed in some logics, and from segmentation and recognition results, higher level interpretations of an image can be derived using such a method of abductive reasoning (Atif/Hudelot/Bloch 2014; Yang/Atif/Bloch 2015). Then the image understanding problem itself is formulated as an explanatory process. The logic is endowed with fuzzy semantics, used to cope with imprecise statements in the knowledge base, such as "the lateral ventricles are dark in T1 weighted magnetic resonance images, the caudate nuclei are external to the lateral ventricles and close to them". Observation is the image and results from segmentation and recognition procedures. Hence, there is an interpretation on two levels: first at the object level, using the approaches presented in the previous sections involving fuzzy representations and structural models, and secondly globally, at the scene level. The advantages of using abstract formulation in a logic is that this second, higher level, interpretation can

3 Note that this is very natural, and explored since the antiquity, while it is much more difficult with machine learning that performs mostly inductions.

take intelligible forms, such as "this image presents an enhanced tumor, which is subcortical and has a small deforming impact on the other structures".

The language in which the knowledge is expressed should be defined according to the granularity level expected of the interpretation and based on whom the description is dedicated to (the explainee). For instance, the description of the content of a pathological brain image will depend on whether the explainee is anyone (without assuming any particular expertise), the patient, or a medical expert who wants to make a decision guided by this description and aims to interact with other experts. Other important questions are related to what should be explained. For instance, a medical expert needs mostly explanations of a result rather than explanations of every step of the algorithm as well as explanations of the links between the results, the data, and the available knowledge. More importantly, explanations are required when the results are unexpected. This is related to the question of when an explanation is needed and refers to the idea of contrastive explanations (why is the result A, when B was expected?).

To go further, another level of explanation is to identify which part of the knowledge base has actually been involved in the reasoning process or is relevant in the object or scene description. An implicit method to do so was mentioned above (Riva et al. 2022). More explicit methods are also very relevant for providing meaningful explanations to users. Fuzzy sets are then useful for establishing a link between the results derived from the image and concepts expressed in the knowledge base, as mentioned at the end of section 2. A simple example is to assess to which degree a spatial relation is satisfied between the resulting objects. Then explanations such as "this object is the left caudate nucleus because it is close to the left ventricle and to the left of it" are easy to derive. For instance, a given spatial relation between two identified objects can be computed, as a number or as a distribution, and then compared to the fuzzy model of this relation (Bloch/Atif 2016). An approach based on fuzzy frequent itemset mining has also been proposed (Pierrard/Poli/Hudelot 2021). Considering the example of structure recognition based on spatial reasoning, explanations become natural by identifying the spatial relations that actually play a role in the recognition. Furthermore, we can make use of hedges and quantifiers to find out whether "most" of the relations in a given set are indeed satisfied by a result, or involved in the image understanding process.

In all that precedes, hybrid AI and the combination of several approaches are at the core of:

- knowledge representation (object properties and relations between objects),
- associating attribute definition with structural frameworks such as graphs hypergraphs or other computational models representing the structure (in the sense of spatial organization) of a scene,
- semantics of logics,
- semantic gap solving,
- spatial reasoning for image understanding, computing similarities between a model and a result,
- providing descriptions of an image in a given language, providing cues for explainability.

They are the main medium to travel from knowledge to data and conversely explain results obtained from data according to the available knowledge.

6. Discussion

To go further in the field of hybrid AI and XAI for image understanding, principles expressed and discussed more generally in AI could be instantiated in this particular domain of application and pave the way for new research directions.

This starts with the definition of interpretability and explainability. An interesting distinction is proposed by Denis and Varenne (2022), where interpretability is defined as the composition of elements that are meaningful for humans, while explanation is strongly related to causality, and understanding is linked to unifying diversity under a common principle (this is may be somewhat different when interpreting an individual image as in medical imaging). In the works summarized in this paper, fuzzy sets are an example that can be used to make explicit the components of knowledge and image information that are involved in a reasoning process. This is done in a semi-qualitative way, close to human understanding, and therefore directly useful to provide explanations.

Seeing explanations as causality has been widely addressed, in particular by Halpern and Pearl (Halpern/Pearl 2005a; Halpern/Pearl 2005b) and by Miller (Miller 2019; Miller 2021), where structural models play a major role. Links with argumentation frameworks (Munro et al. 2022) and extensions of contrastive explanations for fuzzy sets (Bloch/Lesot 2022) have recently been proposed. Notions such as contrast and relevance are put to the fore, and would be also important to consider in image understanding. For instance, explaining why a certain decision was proposed by an algorithm, and not another, is a way to make explanations more convincing. A simple way to do so based on the methods presented here would be to compare resulting image descriptions with different models or decisions, and to identify which components in the knowledge or in the reasoning was responsible for a particular decision proposal. This would be particularly interesting in medical imaging, where explanations are mostly required when the result provided by an algorithm differs from the expected one. This deserves further investigation. The level of explanation should depend on the explainee, as mentioned above, and a deeper study of this aspect could take inspiration from the work on intelligibility by Coste-Marquis and Marquis (2020) (for instance based on projections on a given vocabulary). This goes with the idea of a human-centered evaluation of AI systems.

It has been advocated by Marcus (2020: 1) that new research should aim at developing "a hybrid, knowledge driven, reasoning based approach, centered around cognitive models, that could provide the substrate for a richer, more robust AI than is currently possible." This is exactly what research in image understanding based on hybrid AI is trying to do, but still at a modest level. The question of bias is related to the one of robustness. Statistical biases, on the one hand, are usually quite well identified in medical imaging. They may come from the limited data, from the under-representativity of parts of a population, from the specificities of the study (which intrinsically limit the population) and of the imaging center to the evolution of the data and the update of the algorithms, etc. This raises difficulties to adapt a method to a different population for instance. One may also wonder whether learning methods implicitly use information that can be relevant or that can be biased (which is then not explicitly identified). On the other hand, cognitive biases (such as confirmation, framing, complacency biases) may be more difficult to assess. An interesting direction of research is to investigate how hybrid AI can cope with these questions.

Finally, it would be interesting to investigate more deeply to which extent hybrid AI and XAI could help answering questions related to ethics, for instance in radiology, where these questions are often raised.

Acknowledgements

The author would like to thank all her co-authors and emphasize that the ideas summarized in this paper benefitted from many joint works with PhD candidates, post-doctoral researchers, colleagues in universities, research centers in several countries, as well as university hospitals and industrial partners. This work was partly supported by the author's chair in Artificial Intelligence (Sorbonne Université and SCAI). A part of the work was performed while the author was with LTCI, Télécom Paris, Institut Polytechnique de Paris.

List of references

Aiello, Marco/Pratt-Hartmann, Ian/Benthem, Johan (eds.) (2007): Handbook of Spatial Logics, Dordrecht: Springer.

Aiguier, Marc/Atif, Jamal/Bloch, Isabelle/Pino Pérez, Ramón (2018): "Explanatory Relations in Arbitrary Logics Based on Satisfaction Systems, Cutting, and Retraction." In: International Journal of Approximate Reasoning 102, pp. 1–20.

Aiguier, Marc/Bloch, Isabelle (2019): "Logical Dual Concepts Based on Mathematical Morphology in Stratified Institutions." In: Journal of Applied Non-Classical Logics 29/4, pp. 392–429.

Aldea, Emanuel/Bloch, Isabelle (2010): "Toward a Better Integration of Spatial Relations in Learning with Graphical Models." In: Fabrice Guillet/Gilbert Ritschard/Djamel Abdelkader Zighed/Henri Briand (eds.), Advances in Knowledge Discovery and Management, Berlin and Heidelberg: Springer, pp. 77–94.

Atif, Jamal/Hudelot, Céline/Bloch, Isabelle (2014): "Explanatory Reasoning for Image Understanding Using Formal Concept Analysis and Description Logics." In: IEEE Transactions on Systems, Man and Cybernetics: Systems 44/5, pp. 552–570.

Atif, Jamal/Hudelot, Céline/Fouquier, Geoffroy/Bloch, Isabelle/Angelini, Elsa (2007): "From Generic Knowledge to Specific Reasoning for Medical Image Interpretation using Graph-based Representations." In: International Joint Conference on Artificial Intelligence IJCAI'07, Hyderabad, India, pp. 224–229.

Avila Garcez, Artur d'/Lamb, Luís C. (2023): "Neurosymbolic AI: the 3rd wave." In: Artificial Intelligence Review (https://doi.org/10.1007/s10462-023-10448-w).

Bloch, Isabelle (1996): "Information Combination Operators for Data Fusion: A Comparative Review with Classification." In: IEEE Transactions on Systems, Man, and Cybernetics 26/1, pp. 52–67.

Bloch, Isabelle (2003): "On Fuzzy Spatial Distances." In: Peter W. Hawkes (ed.), Advances in Imaging and Electron Physics, Amsterdam: Elsevier, pp. 51–122.

Bloch, Isabelle (2005): "Fuzzy Spatial Relationships for Image Processing and Interpretation: A Review." In: Image and Vision Computing 23/2, pp. 89–110.

Bloch, Isabelle (2006): "Spatial Reasoning under Imprecision using Fuzzy Set Theory, Formal Logics and Mathematical Morphology." In: International Journal of Approximate Reasoning 41/2, pp. 77–95.

Bloch, Isabelle (2021a): "Mathematical Morphology and Spatial Reasoning: Fuzzy and Bipolar Setting." In: TWMS Journal of Pure and Applied Mathematics, special issue 12/1, pp. 104–125.

Bloch, Isabelle (2021b): "Modeling Imprecise and Bipolar Algebraic and Topological Relations Using Morphological Dilations." In: Mathematical Morphology – Theory and Applications 5/1, pp. 1–20.

Bloch, Isabelle (2022): "Hybrid Artificial Intelligence for Knowledge Representation and Model-Based Medical Image Understanding – Towards Explainability." In: Discrete Geometry and Mathematical Morphology, Second International Joint Conference, DGMM 2022, Straßbourg, France, pp. 17–25.

Bloch, Isabelle/Atif, Jamal (2016): "Defining and Computing Hausdorff Distances Between Distributions on the Real Line and on the Circle: Link Between Optimal Transport and Morphological Dilations." In: Mathematical Morphology: Theory and Applications 1/1, pp. 79–99.

Bloch, Isabelle/Géraud, Thierry/Maître, Henri (2003): "Representation and Fusion of Heterogeneous Fuzzy Information in the 3D Space for Model-Based Structural Recognition – Application to 3D Brain Imaging." In: Artificial Intelligence 148, pp. 141–175.

Bloch, Isabelle/Lang, Jérôme/Pino Pérez, Ramón/Uzcátegui, Carlos (2018): Morphology for Knowledge Dynamics: Revision, Fusion, Abduction, arXiv Preprint (https://arxiv.org/abs/1802.05142).

Bloch, Isabelle/Lesot, Marie-Jeanne (2022): "Towards a Formulation of Fuzzy Contrastive Explanations." In: IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), Padua, Italy, pp. 1–8.

Bloch, Isabelle/Ralescu, Anca (2023): Fuzzy Sets Methods in Image Processing and Understanding: Medical Imaging Applications, Cham: Springer Nature.

Cesar, Roberto M. Jr./Bengoetxea, Endika/Bloch, Isabelle/Larranaga, Pedro (2005): "Inexact Graph Matching for Model-Based Recognition: Evaluation and Comparison of Optimization Algorithms." In: Pattern Recognition 38/11, pp. 2099–2113.

Chopin, Jérémy/Fasquel, Jean-Baptiste/Mouchère, Harold/Dahyot, Rozenn/Bloch, Isabelle (2022): "Improving Semantic Segmentation with Graph-Based Structural Knowledge." In: International Conference on Pattern Recognition and Artificial Intelligence (ICPRAI 2022), Paris, France, pp. 173–184.

Colliot, Olivier/Camara, Oscar/Bloch, Isabelle (2006): "Integration of Fuzzy Spatial Relations in Deformable Models – Application to Brain MRI Segmentation." In: Pattern Recognition 39/8, pp. 1401–1414.

Coradeschi, Silvia/Saffiotti, Alessandro (1999): "Anchoring Symbols to Vision Data by Fuzzy Logic." In: European Conference on Symbolic and Quantitative Approaches with Uncertainty (ECSQARU'99), London, UK, pp. 104–115.

Coste-Marquis, Sylvie/Marquis, Pierre (2020): "From Explanations to Intelligible Explanations." In: 1st International Workshop on Explainable Logic-Based Knowledge Representation (XLoKR'20), 2020, Rhodes, Greece (https://univ-artois.hal.science/hal-03300380/document).

Couteaux, Vincent/Si-Mohamed, Salim/Nempont, Olivier/Lefevre, Thierry/Popoff, Alexandre/Pizaine, Guillaume/Villain, Nicolas/Bloch, Isabelle/Cotten, Anne/Boussel, Loic (2019): "Automatic Knee Meniscus Tear Detection and Orientation Classification with Mask-RCNN." In: Diagnostic and Interventional Imaging 100/4, pp. 235–242.

De Raedt, Luc/Dumančić, Sebastijan/Manhaeve, Robin/Marra, Giuseppe (2020): "From Statistical Relational to Neuro-Symbolic Artificial Intelligence." In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20), Yokohama, Japan, pp. 4943–4950.

Delmonte, Alessandro/Mercier, Corentin/Pallud, Johan/Bloch, Isabelle/Gori, Pietro (2019): "White Matter Multi-Resolution Segmentation Using Fuzzy Set Theory." In: IEEE International Symposium on Biomedical Imaging (ISBI), Venice, Italy, pp. 459–462.

Denis, Cristophe/Varenne, Franck (2022): "Interprétabilité et explicabilité de phénomènes prédits par de l'apprentissage machine." In: Revue Ouverte d'Intelligence Artificielle 3/3-4, pp. 287–310.

Deruyver, Aline/Hodé, Yann (1997): "Constraint Satisfaction Problem with Bilevel Constraint: Application to Interpretation of Over-Segmented Images." In: Artificial Intelligence 93/1-2, pp. 321–335.

Dubois, Didier/Prade, Henri (1985): "A Review of Fuzzy Set Aggregation Connectives." In: Information Sciences 36/1-2, pp. 85–121.

Fasquel, Jean-Baptiste/Delanoue, Nicolas (2019): "A Graph-Based Image Interpretation Method Using A Priori Qualitative Inclusion and Photometric Relationships." In: IEEE Transactions on Pattern Analysis and Machine Intelligence 41/5, pp. 1043–1055.

Fouquier, Geoffroy/Atif, Jamal/Bloch, Isabelle (2012): "Sequential Model-Based Segmentation and Recognition of Image Structures Driven by Visual Features and Spatial Relations." In: Computer Vision and Image Understanding 116/1, pp. 146–165.

Freeman, John (1975): "The Modelling of Spatial Relations." In: Computer Graphics and Image Processing 4/2, pp. 156–171.

Garnelo, Marta/Shanahan, Murray (2019): "Reconciling Deep Learning with Symbolic Artificial Intelligence: Representing Objects and Relations." In: Current Opinion in Behavioral Sciences 29, pp. 17–23.

Halpern, Joseph Y./Pearl, Judea (2005a): "Causes and Explanations: A Structural-Model Approach. Part I: Causes." In: The British Journal for the Philosophy of Science 56/4, pp. 843–887.

Halpern, Joseph Y./Pearl, Judea (2005b): "Causes and Explanations: A Structural-Model Approach. Part II: Explanations." In: The British Journal for the Philosophy of Science 56/4, pp. 889–911.

Harnad, Stevan (1990): "The symbol grounding problem." In: Physica 42/1-3, pp. 335–346.

Hudelot, Céline/Atif, Jamal/Bloch, Isabelle (2008): "Fuzzy Spatial Relation Ontology for Image Interpretation." In: Fuzzy Sets and Systems 159, pp. 1929–1951.

Kahneman, Daniel (2012): Thinking, Fast and Slow, New York: Penguin.

Kautz, Henry A. (2022): "The third AI summer: AAAI Robert S. Engelmore Memorial Lecture." In: AI Magazine 43/1, pp. 93–104.

Landini, Gabriel/Galton, Antony/Randell, David/Fouad, Shereen (2019): "Novel Applications of Discrete Mereotopology to Mathematical Morphology." In: Signal Processing: Image Communication 76, pp. 109–117.

Marcus, Gary (2020): The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence, arXiv Preprint (https://arxiv.org/abs/2002.06177).
Miller, Tim (2019): "Explanation in Artificial Intelligence: Insights from the Social Sciences." In: Artificial Intelligence 267, pp. 1–38.
Miller, Tim (2021): "Contrastive Explanation: A Structural-Model Approach." In: The Knowledge Engineering Review 36, e14.
Munro, Yann/Bloch, Isabelle/Chetouani, Mohamed/Lesot, Marie-Jeanne/Pelachaud, Catherine (2022): "Argumentation and Causal Models in Human-Machine Interaction: A Round Trip." In: 8th International Workshop on Artificial Intelligence and Cognition, Örebro, Sweden (https://hal.sorbonne-universite.fr/hal-03739310).
Nempont, Olivier/Atif, Jamal/Bloch, Isabelle (2013): "A Constraint Propagation Approach to Structural Model-Based Image Segmentation and Recognition." In: Information Sciences 246, pp. 1–27.
Perchant, Aymeric/Bloch, Isabelle (2002): "Fuzzy Morphisms between Graphs." In: Fuzzy Sets and Systems 128/2, pp. 149–168.
Pierrard, Régis/Poli, Jean-Philippe/Hudelot, Céline (2021): "Spatial Relation Learning for Explainable Image Classification and Annotation in Critical Applications." In: Artificial Intelligence 292, 103434.
Randell, David A./Cui, Zhan/Cohn, Anthony G. (1992): "A Spatial Logic based on Regions and Connection." In: Principles of Knowledge Representation and Reasoning. Proceedings of the Third International Conference, Cambridge, MA, USA, pp. 165–176.
Riva, Mateus/Gori, Pietro/Yger, Florian/Bloch, Isabelle (2022): "Is the U-Net Directional-Relationship Aware?" In: 2022 IEEE International Conference on Image Processing (ICIP), Bordeaux, France, pp. 3391–3395.
Schockaert, Steven/De Cock, Martine/Cornelis, Chris/Kerre, Etienne E. (2008): "Fuzzy Region Connection Calculus: Representing Vague Topological Information." In: International Journal of Approximate Reasoning 48/1, pp. 314–331.
Schockaert, Steven/De Cock, Martine/Kerre, Etienne E. (2009): "Spatial Reasoning in a Fuzzy Region Connection Calculus." In: Artificial Intelligence 173/2, pp. 258–298.
Smeulders, Arnold/Worring, Marcel/Santini, Simone/Gupta, Amarnath/Jain, Ramesh (2000): "Content-based Image Retrieval at the End of the Early Years." In: IEEE Transactions on Pattern Analysis and Machine Intelligence 22/12, pp. 1349–1380.

Virzi, Alessio/Gori, Pietro/Muller, Cécile Olivia/Mille, Eva/Peyrot, Quoc/Berteloot, Laureline/Boddaert, Nathalie/Sarnacki, Sabine/Bloch, Isabelle (2018): "Segmentation of Pelvic Vessels in Pediatric MRI Using a Patch-based Deep Learning Approach." In: PIPPI MICCAI Workshop, Granada, Spain, pp. 97–106.

Xie, Xiaozheng/Niu, Jianwei/Liu, Xuefeng/Chen, Zhengsu/Tang, Shaojie/Yu, Shui (2021): "A Survey on Incorporating Domain Knowledge into Deep Learning for Medical Image Analysis." In: Medical Image Analysis 69, 101985.

Yager, Ronald R. (1991): "Connectives and Quantifiers in Fuzzy Sets." In: Fuzzy Sets and Systems 40, pp. 39–75.

Yang, Y./Atif, Jamal/Bloch, Isabelle (2015): "Abductive Reasoning Using Tableau Methods for High-level Image Interpretation." In: KI 2015: Advances in Artificial Intelligence, 38th Annual German Conference on AI, Dresden, Germany, pp. 356–365.

Zadeh, Lotfi A. (1965): "Fuzzy Sets." In: Information and Control 8/3, pp. 338–353.

Zadeh, Lotfi A. (1975): "The Concept of a Linguistic Variable and its Application to Approximate Reasoning." In: Information Sciences 8/3, pp. 199–249.