Journal of Computational Intelligence and Decision Analytics

Journal homepage: www.cida-journal.org

ISSN: 3102-6125

The Principle of Bounded Fuzziness for Uncertainty in Fuzzy Sets

Tamunokuro Opubo William-West[1,*], Paul Augustine Ejegwa[2]

[1] Department of Mathematics, Faculty of Physical Sciences, Ahmadu Bello University, Zaria, Nigeria
[2] Department of Mathematics, Federal University of Health Sciences, Otukpo, Nigeria

ARTICLE INFO

Article history:
Received 24 April 2026
Received in revised form 29 May 2026
Accepted 10 August 2026
Available online 15 August 2026

Keywords:
Fuzzy sets; Principle of Bounded Fuzziness; Ball Fuzzy Sets; Epistemic uncertainty.

ABSTRACT

The representation of epistemic uncertainty is a fundamental issue in fuzzy set theory and approximate reasoning. Existing frameworks accommodate multiple admissible assessments but provide little insight into the structural form that epistemic uncertainty should assume. This paper introduces the *Principle of Bounded Fuzziness* (PBF), which postulates that admissible assessments generated under shared knowledge remain confined to a bounded neighborhood around a representative assessment in an appropriate metric space. Specializing the PBF to the space of membership functions $[0,1]^X$ yields *Ball Fuzzy Sets*, defined as Chebyshev metric balls that admit a canonical three-way decomposition into belief, epistemic uncertainty, and disbelief satisfying the conservation relation $b(x) + u(x) + d(x) = 1$. We develop the structural and algebraic foundations of ball fuzzy sets, establish their principal geometric properties, derive complement, intersection, and union operations, and examine their relationships with interval-valued, hesitant, and higher-order fuzzy representations. Empirical studies on human probability distributions and expert medical assessments provide evidence that bounded disagreement is an observable phenomenon. In particular, finite coverage radii exist at multiple confidence levels; assessments derived from shared knowledge exhibit significantly tighter clustering than random baselines ($p < 0.001$); the estimated bounds remain stable under resampling; and large deviations become increasingly unlikely. These observations are consistent with the central premise that shared knowledge constrains admissible disagreement. An ICU risk assessment study illustrates practical utility of the proposed framework for uncertainty-aware reasoning. By providing a geometric interpretation of epistemic uncertainty, the Principle of Bounded Fuzziness complements existing principles of uncertainty invariance and justifiable granularity.

1. Introduction

The introduction of fuzzy sets by Zadeh [1] provided a powerful framework for representing concepts whose boundaries are inherently gradual rather than crisp. By replacing binary membership with a membership function

$$\mu : X \to [0,1],$$

fuzzy set theory enables the representation of partial belonging and has become a foundational component of intelligent systems, pattern recognition, control, artificial intelligence and decision analysis. [2–5]. The success of the framework stems from a simple insight: many concepts encountered in practice are not naturally dichotomous and therefore require graded rather than binary representation.

While the classical framework successfully models gradualness, practical applications often reveal an additional layer of uncertainty. In many situations, the uncertainty does not arise solely from the concept itself but from the assessment of its membership grade. Different experts may assign different evaluations to the same patient, different sensors may produce slightly different measurements of the same

*Corrsponding author.
E-mail address: westtamuno@gmail.com

quantity, and different predictive models may produce different outputs for the same observation. Consequently, uncertainty may exist not only in the concept being represented but also in the numerical assessment used to represent it.

This distinction is fundamental. Classical fuzziness concerns the graded nature of a concept. Epistemic uncertainty concerns the reliability of the assigned grade. A patient may belong to the class "high risk" with degree 0.85, yet different physicians may reasonably assign values such as 0.78, 0.83, or 0.90. The first phenomenon concerns gradual membership. The second concerns uncertainty about the membership assessment itself.

Numerous extensions of fuzzy sets have been developed to address such situations. Interval-valued fuzzy sets [6], hesitant fuzzy sets [7], type-2 fuzzy sets [8], intuitionistic fuzzy sets, and related models allow multiple admissible assessments to coexist. These frameworks have significantly enriched the expressive power of uncertainty modeling. Nevertheless, a fundamental question remains largely unexplored.

> *If multiple admissible assessments exist, what structural form should their uncertainty possess?*

This question is not specific to fuzzy sets. It arises whenever knowledge is incomplete and assessments are produced by human experts, measurement devices, or computational models. Consider a physician estimating patient risk, an engineer evaluating system reliability, or a machine learning model predicting class membership. In each case, multiple admissible assessments may arise. Yet the disagreement among these assessments is rarely arbitrary.

The key observation motivating this work is that disagreement is typically constrained by the knowledge available to the assessors. Physicians trained on the same medical evidence may disagree, but their disagreement is limited by shared clinical knowledge. Calibrated sensors may produce different measurements, but the variation is constrained by the physical properties of the sensing process. Well-trained predictive models may produce different outputs, but their predictions remain influenced by the same underlying data-generating mechanism.

The existence of disagreement therefore does not imply unrestricted variation. On the contrary, the presence of shared evidence, common knowledge, or common observational mechanisms tends to constrain the range of admissible assessments.

This observation suggests a more fundamental principle:

> *Shared knowledge induces constrained disagreement.*

If disagreement is constrained, then admissible assessments cannot diverge arbitrarily from a representative estimate. Instead, they must remain confined to a finite neighborhood surrounding that estimate. The central issue is therefore not whether uncertainty exists, but how the geometry of admissible uncertainty should be represented.

1.1 The Need for a Structural Principle

The development of uncertainty theories has often been guided by foundational principles. Klir's Principle of Uncertainty Invariance [9] emphasizes consistency across equivalent uncertainty representations. Pedrycz's Principle of Justifiable Granularity [10] provides a methodology for constructing information granules by balancing coverage and specificity. Such principles do not merely define mathematical objects; they explain why particular structures should arise.

However, existing uncertainty theories provide relatively little guidance regarding the geometric organization of admissible assessments. Most frameworks permit multiple possible evaluations, but they rarely explain why certain evaluations should be regarded as admissible while others should not.

Consequently, an important structural question remains:

> *Once uncertainty is acknowledged to exist, what geometric form should admissible uncertainty assume within a representation space?*

We argue that the answer follows naturally from the existence of constrained disagreement. If assessments are constrained by shared knowledge, then admissible assessments must remain within a finite distance of a representative assessment. The resulting uncertainty is therefore naturally represented by a bounded region.

1.2 The Principle of Bounded Fuzziness

The foregoing observations motivate the Principle of Bounded Fuzziness (PBF). Rather than viewing uncertainty as an arbitrary collection of admissible assessments, the PBF regards uncertainty as a bounded neighborhood surrounding a representative assessment.

Let $\mathcal{E}$ denote a space of admissible assessments and let $e_0 \in \mathcal{E}$ denote a representative assessment obtained from the available knowledge. The PBF asserts that admissible assessments remain confined to a bounded region centered at e_0.

Principle 1.1 (Principle of Bounded Fuzziness)**.** Let $\mathcal{E}$ be a space of admissible assessments and let $e_0 \in \mathcal{E}$ be a representative assessment. Epistemic uncertainty shall be represented by a bounded neighborhood

$$B_r(e_0) = \{e \in \mathcal{E} : d(e, e_0) \leq r\},$$

where d is a metric on $\mathcal{E}$ and $r \geq 0$ is a tolerance parameter quantifying the extent of admissible disagreement.

The principle contains two independent components:

1. A *reference assessment* e_0, representing the best available estimate or consensus assessment.
2. A *tolerance parameter* r, representing the extent to which admissible assessments may deviate from the reference.

Together, these quantities characterize both the location and the extent of epistemic uncertainty.

It is important to emphasize that the PBF adopts an epistemic interpretation. The reference assessment need not represent a metaphysical truth. It may instead be interpreted as a consensus estimate, a central tendency, or the best assessment available under the current state of knowledge. The principle concerns the structure of admissible disagreement around that reference.

1.3 From the Principle to Ball Fuzzy Sets

The primary objective of this work is not to postulate a new fuzzy set model but to derive an uncertainty representation from the Principle of Bounded Fuzziness.

To achieve this, we consider the space

$$\mathcal{E} = [0, 1]^X,$$

where each element represents a membership function on a finite universe X.

The PBF requires that admissible membership functions remain uniformly close to a reference membership function. This requirement is naturally expressed through the coordinate-wise condition

$$|\mu(x) - \mu_0(x)| \leq r, \qquad \forall x \in X.$$

Among the standard metrics on $[0, 1]^X$, the Chebyshev metric

$$d_\infty(\mu, \nu) = \max_{x \in X} |\mu(x) - \nu(x)|$$

is the unique standard metric that directly captures this coordinate-wise interpretation. Unlike the L_1 and L_2 metrics, which permit compensation among coordinates, the Chebyshev metric enforces a uniform bound at every element of the universe.

Consequently, the PBF induces the admissible region

$$B_r(\mu_0) = \{\mu \in [0, 1]^X : d_\infty(\mu, \mu_0) \leq r\},$$

which forms a bounded neighborhood around the reference membership function. We refer to this structure as a *ball fuzzy set*.

The resulting representation provides a natural decomposition into belief, uncertainty, and disbelief components and establishes a direct connection between epistemic uncertainty and geometric structure.

1.4 Empirical Motivation

A structural principle is meaningful only if its implications are observable in practice. The central empirical claim underlying the PBF is that disagreement generated under shared knowledge remains bounded.

To examine this claim, we investigate two fundamentally different uncertainty-generation mechanisms. The first is CIFAR-10H, which contains human probability distributions obtained from multiple annotators performing image classification. The second is CURVAS, which represents expert variation in medical image analysis.

Rather than treating these datasets as applications, we use them to examine the observable consequences of the principle itself. In particular, we investigate whether disagreement remains confined within finite probabilistic bounds, whether shared knowledge reduces disagreement relative to random assessments, whether the estimated bounds remain stable under resampling, and whether large deviations become increasingly unlikely.

The results consistently indicate that disagreement generated under shared information is substantially more constrained than disagreement generated under random mechanisms. These findings provide empirical evidence that bounded disagreement is not merely a mathematical assumption but an observable characteristic of epistemic uncertainty.

The main contributions of this work are summarized as follows.

1. A foundational principle. We introduce the Principle of Bounded Fuzziness as a structural doctrine for representing epistemic uncertainty in membership assessment.

2. *A geometric interpretation of uncertainty.* We show that constrained disagreement naturally leads to bounded neighborhoods in a suitable representation space.

3. *Ball fuzzy sets.* We derive ball fuzzy sets as the canonical realization of the Principle of Bounded Fuzziness within the space of membership functions.

4. *Structural and algebraic analysis.* We establish the principal geometric and algebraic properties of the resulting representation.

5. *Empirical examination of the principle.* We investigate multiple observable implications of bounded epistemic disagreement using human and expert assessment data.

6. *Practical demonstration.* We illustrate the applicability of the framework in uncertainty-aware risk assessment.

2. Foundational Principles

The Principle of Bounded Fuzziness (PBF) occupies a distinct position in the landscape of uncertainty modeling. It is neither a meta-theoretical principle like Klir's invariance nor an operational principle like Pedrycz's justifiable granularity. Rather, it addresses the structural question: given a representation space, what geometric form should epistemic uncertainty take? This section develops the PBF from first principles, establishes its formal foundation, and justifies its key components.

2.1 The Principle

The empirical observation that motivates the PBF is simple yet pervasive: in contexts where multiple observers assess the same graded concept, their assessments vary, but this variation is bounded. Physicians assessing patient risk, sensors measuring physical quantities, and classifiers producing probability scores all yield values that cluster within limits determined by shared knowledge and measurement constraints.

This boundedness is not a flaw in the assessment process; it reflects the epistemic structure of the situation. When we model a concept with a membership function, we are not claiming that the function is known with certainty. Rather, we are acknowledging that different reasonable assessments are possible, but only within limits. The PBF makes this structure explicit.

Principle 2.1 (Principle of Bounded Fuzziness). Let $\mathcal{E}$ be a space of admissible objects representing possible assessments of a concept. For a given reference assessment $e_0 \in \mathcal{E}$ and a tolerance $r \geq 0$, epistemic uncertainty about the true assessment is represented by the closed metric neighborhood

$$B_r(e_0) = \{e \in \mathcal{E} \mid d(e, e_0) \leq r\},$$

where d is a metric on $\mathcal{E}$.

The PBF rests on two structural components:

1. *The reference* e_0: the location parameter, representing the best available estimate or consensus assessment. This is what we would commit to if forced to choose a single value.
2. *The tolerance* r: the scale parameter, representing the maximum permissible deviation from the reference. This quantifies the extent of epistemic uncertainty—the degree of doubt about the exact assessment.

These components are structurally independent. The reference determines *where* the admissible region is centered; the tolerance determines *how large* that region is. They may be specified, estimated, and updated independently, reflecting the fact that improving our estimate (refining the reference) and gaining confidence (reducing the tolerance) are distinct epistemic activities.

It is important to clarify two points of interpretation. First, the PBF adopts an *epistemic* rather than ontological interpretation. It does not assume the existence of a "true" membership function in any metaphysical sense. Rather, it assumes that for a given concept and context, there is an ideal assessment that would be obtained with perfect information and complete rationality—a notion analogous to the "true value" in measurement theory, which serves as a useful idealization even when we acknowledge that measurements are always subject to error. Readers who prefer a weaker interpretation may regard e_0 simply as a consensus reference point.

Second, the PBF addresses *second-order epistemic uncertainty*—uncertainty about the membership grade itself—not *first-order fuzziness* (the gradedness of the concept). Classical fuzziness concerns the fact that membership is gradual: a patient can be "high risk" to degree 0.85. The PBF concerns the fact that different observers may assign 0.78, 0.85, or 0.92 to the same patient. These are complementary phenomena, and the PBF is not intended to replace the classical interpretation. The term "bounded fuzziness" refers specifically to the boundedness of this second-order uncertainty.

2.2 The Choice Metric

The PBF is metric-agnostic at the level of the general principle. However, when applied specifically to membership functions in $\mathcal{E} = [0, 1]^X$, the choice of metric carries semantic consequences that are not arbitrary.

The empirical motivation for the PBF—bounded disagreement per element—requires that for *every* $x \in X$, the deviation $|\mu(x) - \mu_0(x)|$ does not exceed the tolerance r. This is a *coordinate-wise* bound: each element's membership grade is independently constrained.

Among the standard metrics on $[0, 1]^X$, only the Chebyshev (supremum) metric

$$d_\infty(\mu, \nu) = \max_{x \in X} |\mu(x) - \nu(x)|$$

directly enforces this coordinate-wise interpretation. If $d_\infty(\mu, \nu) \leq r$, then *every* coordinate satisfies $|\mu(x) - \nu(x)| \leq r$. Conversely, if every coordinate satisfies the bound, then $d_\infty(\mu, \nu) \leq r$.

Metrics such as L_1 or L_2 permit trade-offs across coordinates: a large deviation at one coordinate can be compensated by small deviations elsewhere, violating the per-element bound. These metrics are appropriate when uncertainty is measured as an average over the universe, but they do not capture the semantic requirement of uniform per-element tolerance.

Chebyshev Ball vs. Euclidean Ball

Only the Chebyshev metric enforces coordinate-wise per-element bounds

Chebyshev Ball (L_∞) **Euclidean Ball** (L_2)

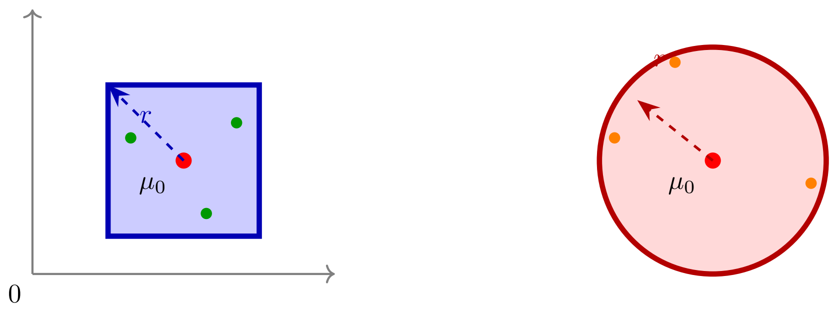

Uniform per-element bound
$|\mu(x) - \mu_0(x)| \le r \ \forall x$

Allows trade-offs across coordinates:
large deviation at one point
compensated by small deviations elsewhere

The following proposition establishes the uniqueness of the Chebyshev metric for this interpretation.

Proposition 2.1. *Let $\mathcal{E} = [0,1]^X$ and let d be a metric satisfying: for all $\mu, \nu \in \mathcal{E}$ and all $r \ge 0$,*

$$d(\mu,\nu) \le r \iff |\mu(x) - \nu(x)| \le r \ \ \forall x \in X.$$

Then $d = d_\infty$.

Proof. The equivalence directly implies $d(\mu,\nu) = \max_x |\mu(x) - \nu(x)|$ by taking $r = d(\mu,\nu)$ in one direction and $r = \max_x |\mu(x) - \nu(x)|$ in the other. □

This uniqueness result establishes that the Chebyshev metric is not merely convenient but *forced* by the semantic interpretation of independent per-element bounds.

2.3 Ball Fuzzy Sets

With the Chebyshev metric, the admissible neighborhood becomes

$$B_r(\mu_0) = \{\mu \in [0,1]^X \mid \max_{x \in X} |\mu(x) - \mu_0(x)| \le r\}.$$

Since membership values must remain in $[0,1]$, this condition yields the coordinate-wise interval representation

$$\mu(x) \in \big[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)\big] \quad \forall x \in X.$$

Equivalently,

$$B_r(\mu_0) = \prod_{x \in X} \big[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)\big].$$

Thus the admissible set forms an axis-aligned hypercube in $[0,1]^X$. We term this geometric object a *ball fuzzy set*.

The interval representation is not assumed; it emerges as the coordinate-wise projection of the underlying geometric structure. The extremal envelopes

$$\underline{\mu}(x) = \max(0, \mu_0(x) - r), \qquad \overline{\mu}(x) = \min(1, \mu_0(x) + r)$$

provide complete information about the admissible set and give rise to a natural three-way decomposition of the epistemic state at each x:

$$\underbrace{\underline{\mu}(x)}_{\text{belief}} + \underbrace{\overline{\mu}(x) - \underline{\mu}(x)}_{\text{epistemic uncertainty}} + \underbrace{1 - \overline{\mu}(x)}_{\text{disbelief}} = 1.$$

Here, belief is the guaranteed degree of membership, disbelief is the guaranteed degree of non-membership, and epistemic uncertainty is the range of possible variation in the membership grade. This decomposition distinguishes what is known from what remains indeterminate.

2.4 *Parameter Determination*

The parameters μ_0 and r are not arbitrary. They may be determined from the application context through various principled means:

- *Empirical data [5, 11]:* Given expert assessments $\{\mu^{(1)}, \ldots, \mu^{(N)}\}$, set $\mu_0(x) = \frac{1}{N}\sum_i \mu^{(i)}(x)$ and $r = \max_i \|\mu^{(i)} - \mu_0\|_\infty$ (or a percentile).
- *Expert consensus [12]:* μ_0 may be a negotiated reference, r the agreed tolerance.
- *Measurement reliability:* r may be derived from known instrument error characteristics.
- *Operational principles:* Pedrycz's Principle of Justifiable Granularity [10] provides a systematic methodology for determining information granules from data, balancing coverage and specificity. The PBF complements this principle by specifying the geometric form the granule should take.

2.5 *Relation to Existing Frameworks*

The PBF and ball fuzzy sets connect to established frameworks in several ways.

Type-2 fuzzy sets. Type-2 fuzzy sets [13] represent uncertainty pointwise: for each x, the membership grade is itself a fuzzy set over $[0, 1]$. The PBF imposes a global constraint: uncertainty at all x arises from a single tolerance r applied to a reference function μ_0. This global structure enables algebraic operations that preserve the representation type.

Interval-valued fuzzy sets. Ball fuzzy sets induce intervals at each x, but not every interval-valued fuzzy set arises from a ball. The intervals must have uniform length (except at boundaries). Ball fuzzy sets are precisely *uniform-radius* interval-valued fuzzy sets.

Hesitant fuzzy sets. Ball fuzzy sets generate continuous intervals of possible membership grades at each x. Discrete hesitant fuzzy sets [7] can be seen as samples from these intervals, providing a geometric foundation for hesitation.

Granular computing. The ball $B_r(\mu_0)$ is an information granule in the space of membership functions, with center μ_0 representing the prototype and radius r representing tolerance. This positions ball fuzzy sets within the granular computing paradigm [14].

Summary

The PBF provides a structural principle for representing epistemic uncertainty in membership functions. It rests on two independent components—reference (location) and tolerance (scale)—and uniquely leads to the Chebyshev metric when interpreted as coordinate-wise boundedness. The resulting ball fuzzy sets offer a geometric representation with a natural three-way decomposition into belief, epistemic uncertainty, and disbelief. The framework connects to existing uncertainty models while providing a global structural constraint that enables coherent algebraic operations.

The following sections develop the mathematical consequences of this foundation, establishing structural properties (Section 3) and algebraic operations (Section 4), before demonstrating the framework in application (Section 5).

3. Ball Fuzzy Sets and Structural Properties

The Principle of Bounded Fuzziness, formalized in Section 2, asserts that epistemic uncertainty about a membership function should be represented by a metric ball $B_r(\mu_0)$ in the space $[0,1]^X$. This section develops the mathematical consequences of this assertion. We define ball fuzzy sets as the canonical geometric realization of the PBF, derive the natural three-way decomposition of the epistemic state, establish the fundamental axioms and structural properties, and characterize the relationship to interval-valued fuzzy sets.

3.1 Ball Fuzzy Sets: Definition and Geometric Realization

We specialize the PBF to the space of membership functions $\mathcal{E} = [0,1]^X$, where $X = \{x_1, \ldots, x_n\}$ is a finite universe. The Chebyshev metric, justified in Section 2.3 as the unique metric capturing coordinate-wise boundedness, induces the admissible neighborhood

$$B_r(\mu_0) = \{\mu \in [0,1]^X \mid \max_{x \in X} |\mu(x) - \mu_0(x)| \le r\}.$$

Since membership values must remain in the unit interval, this condition is equivalent to the coordinate-wise interval representation

$$\mu(x) \in \big[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)\big] \quad \forall x \in X.$$

This yields the product structure

$$B_r(\mu_0) = \prod_{x \in X} \big[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)\big],$$

which is an axis-aligned hypercube in $[0,1]^X$.

Definition 3.1 (Ball Fuzzy Set)**.** Let $\mu_0 : X \to [0,1]$ be a reference membership function and $r \ge 0$ a global tolerance parameter. The *ball fuzzy set* $\mathcal{B} = B_r(\mu_0)$ is the Chebyshev ball

$$\mathcal{B} = \{\mu \in [0,1]^X \mid \|\mu - \mu_0\|_\infty \le r\}.$$

For each $x \in X$, the coordinate projection is

$$\mathcal{B}(x) = \big[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)\big].$$

Several observations about this definition merit emphasis.

First, the interval-valued representation is *derived* from the geometric construction, not assumed. The fundamental object is the metric ball; the interval representation is its coordinate projection. This reverses the usual approach in interval-valued fuzzy set theory, where intervals are postulated and geometry is secondary.

Second, the representation is *canonical* in the sense that the mapping $B_r(\mu_0) \mapsto (\underline{\mu}, \overline{\mu})$ is injective: distinct balls yield distinct envelope pairs, and the envelope pair determines the ball uniquely. We term $(\underline{\mu}, \overline{\mu})$ the *canonical envelope* of the ball.

Definition 3.2 (Canonical Envelope)**.** For a ball fuzzy set $\mathcal{B} = B_r(\mu_0)$, the *canonical envelope* is the pair of functions $(\underline{\mu}, \overline{\mu})$ defined pointwise by

$$\underline{\mu}(x) = \max(0, \mu_0(x) - r), \qquad \overline{\mu}(x) = \min(1, \mu_0(x) + r).$$

The canonical envelope satisfies $\underline{\mu}(x) \le \overline{\mu}(x)$ for all x, and the ball is recovered as

$$B_r(\mu_0) = \{\mu \in [0,1]^X \mid \underline{\mu}(x) \le \mu(x) \le \overline{\mu}(x)\ \forall x \in X\}.$$

3.2 The Three-Way Decomposition: Belief, Epistemic Uncertainty, and Disbelief

The canonical envelope induces a natural decomposition of the epistemic state at each element $x \in X$. Define:

$$b(x) = \underline{\mu}(x), \qquad u(x) = \overline{\mu}(x) - \underline{\mu}(x), \qquad d(x) = 1 - \overline{\mu}(x).$$

These quantities satisfy

$$b(x) + u(x) + d(x) = 1 \quad \forall x \in X,$$

and admit a direct epistemic interpretation:

- *Belief* $b(x)$: the guaranteed degree of membership. This is the lowest membership grade that any admissible function assigns to x. It represents what we know with certainty.
- *Disbelief* $d(x)$: the guaranteed degree of non-membership. This is the complement of the highest membership grade that any admissible function assigns to x. It represents what we can rule out with certainty.
- *Epistemic uncertainty* $u(x)$: the width of the admissible interval. This quantifies the range of possible variation in the membership grade at x. It represents what remains indeterminate—the extent to which we are uncertain about the exact grade.

Remark 3.1 (Distinction from first-order fuzziness)**.** It is crucial to distinguish $u(x)$ from *first-order fuzziness*. The quantity $u(x)$ captures *second-order epistemic uncertainty*—uncertainty about the membership grade itself. It arises from the bounded tolerance r and quantifies how much the grade may vary among reasonable assessments. This is distinct from the *first-order gradedness* of the concept, which is captured by the reference assessment $\mu_0(x)$. In the language of type-2 fuzzy sets, $u(x)$ corresponds to the secondary uncertainty—the spread of possible membership values—not the primary membership grade.

The decomposition is illustrated in Figure 1. For each x, the interval $[\underline{\mu}(x), \overline{\mu}(x)]$ is partitioned into three regions: belief (lower), epistemic uncertainty (middle), and disbelief (upper). This provides a complete, pointwise characterization of the epistemic state.

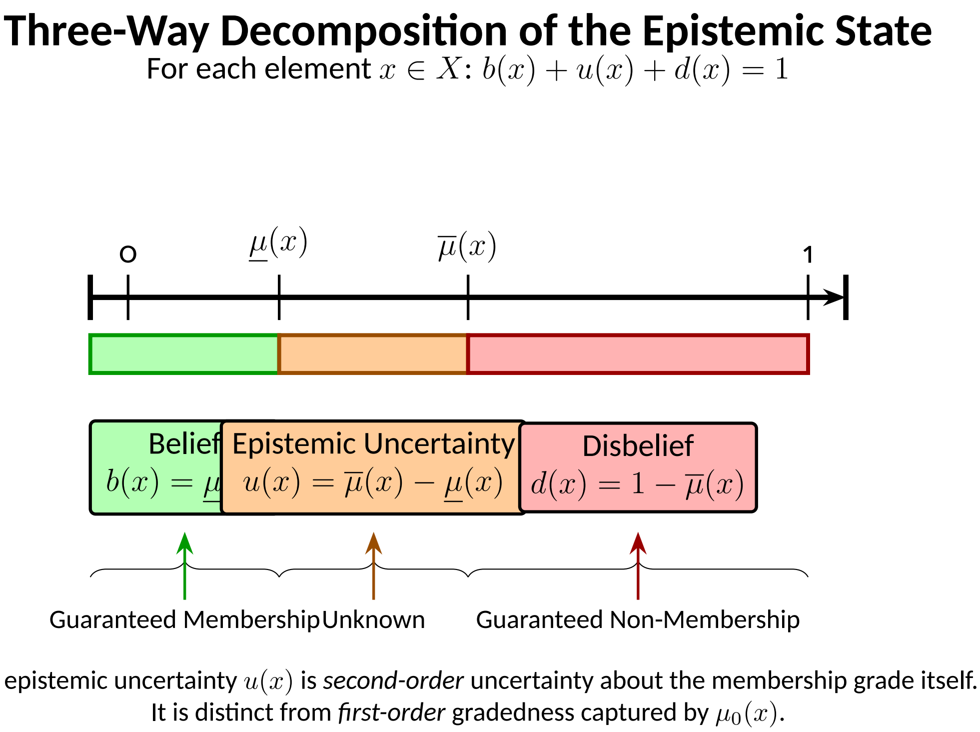

3.3 Fundamental Axioms

The PBF imposes two fundamental axioms on any representation of epistemic uncertainty in a metric space $(\mathcal{E}, d)$.

Axiom 3.1 (Reference and Bound)**.** There exists a reference assessment $e_0 \in \mathcal{E}$ and a tolerance $r \geq 0$ such that all admissible states satisfy $d(e, e_0) \leq r$.

Axiom 3.2 (Closure)**.** The admissible set is closed under the metric topology; it is precisely the closed ball

$$B_r(e_0) = \{e \in \mathcal{E} \mid d(e, e_0) \leq r\}.$$

These axioms capture the essence of bounded fuzziness: uncertainty is quantified by a single global parameter r, and the admissible region is exactly the set of points within that distance of a chosen reference.

From these, several fundamental properties follow directly.

Proposition 3.1 (Monotonicity)**.** *If $r_1 \leq r_2$, then $B_{r_1}(e_0) \subseteq B_{r_2}(e_0)$. Thus uncertainty grows monotonically with the tolerance parameter.*

Corollary 3.1 (Degeneracy)**.** *When $r = 0$, the admissible set reduces to the singleton $\{e_0\}$. Perfect knowledge (zero tolerance) corresponds to a single assessment.*

3.4 Geometric Properties of Ball Fuzzy Sets

When the epistemic space $\mathcal{E}$ is a normed linear space, the axioms yield additional properties that are essential for later developments. These properties establish ball fuzzy sets as mathematically well-behaved objects.

Proposition 3.2 (Convexity)**.** *In a normed linear space, every closed ball $B_r(e_0)$ is convex.*

Proof. For any $x, y \in B_r(e_0)$ and $\lambda \in [0, 1]$,

$$\|\lambda x + (1 - \lambda)y - e_0\| \leq \lambda\|x - e_0\| + (1 - \lambda)\|y - e_0\| \leq r.$$

Thus the weighted average of any two admissible states is also admissible. □

Interpretation 3.1. Convexity ensures that compromise assessments remain within the epistemic tolerance. In multi-expert contexts, averaging admissible opinions yields another admissible opinion.

Proposition 3.3 (Translation Invariance)**.** *In a normed space, $B_r(e_0) = e_0 + B_r(0)$, where $B_r(0) = \{x \mid \|x\| \leq r\}$.*

Interpretation 3.2. Translation invariance reveals a fundamental symmetry: the admissible region around any reference is a translate of the same canonical set. The reference acts as a *location parameter*, while the tolerance is a pure *scale parameter*. This justifies treating the reference and tolerance independently.

Lemma 3.1 (Lipschitz Stability)**.** *For any $e_1, e_2 \in \mathcal{E}$ and any $r \geq 0$,*

$$B_r(e_1) \subseteq B_{r+\|e_1-e_2\|}(e_2).$$

Proof. If $x \in B_r(e_1)$, then $\|x - e_2\| \leq \|x - e_1\| + \|e_1 - e_2\| \leq r + \|e_1 - e_2\|$. □

Interpretation 3.3. Small errors in estimating the reference cause only small expansions of the admissible set. This provides a foundation for sensitivity analysis and stable inference.

Proposition 3.4 (Diameter)**.** *In a normed linear space, $\operatorname{diam}(B_r(e_0)) = 2r$.*

Proof. For any $x, y \in B_r(e_0)$, $\|x - y\| \leq 2r$; equality is achieved by points on opposite sides of the center. □

Interpretation 3.4. The diameter $2r$ gives a global measure of total epistemic uncertainty: it bounds the maximum possible disagreement between any two admissible membership functions. This directly links the geometric parameter r to the intuitive notion of ”spread” in expert opinions.

Theorem 3.1 (Compactness). *If $\mathcal{E}$ is finite-dimensional, every closed ball $B_r(e_0)$ is compact.*

Proof. Finite-dimensional normed spaces are locally compact; closed balls are closed and bounded, hence compact. □

Interpretation 3.5. Compactness guarantees that continuous functions (e.g., expected utility, risk measures) attain their maxima and minima over the admissible set. This is crucial for decision-making under uncertainty: it ensures that optimal choices exist within the ball.

Theorem 3.2 (Center Identifiability). *Let $\mathcal{E}$ be a normed linear space. If $B_r(e_0) = B_r(e_1)$ for some $r > 0$, then $e_0 = e_1$.*

Proof. Assume $B_r(e_0) = B_r(e_1)$. Suppose $e_0 \neq e_1$. Define $u = \frac{e_0 - e_1}{\|e_0 - e_1\|}$ and consider $x = e_0 + ru$. Then $\|x - e_0\| = r$, so $x \in B_r(e_0)$. But

$$\|x - e_1\| = \|(e_0 - e_1) + ru\| = \|e_0 - e_1\| + r > r,$$

so $x \notin B_r(e_1)$. Contradiction. Therefore $e_0 = e_1$. □

Interpretation 3.6. The ball uniquely determines its center. There is no ambiguity about which reference generated the admissible set. This justifies treating μ_0 as a meaningful epistemic summary.

Theorem 3.3 (Intersection Condition). *Let $e_0, e_1 \in \mathcal{E}$ and $r, s \geq 0$.*

(i) If $B_r(e_0) \cap B_s(e_1) \neq \emptyset$, then $d(e_0, e_1) \leq r + s$.

(ii) Conversely, if $\mathcal{E}$ is a normed linear space and $d(e_0, e_1) \leq r + s$, then the intersection is non-empty.

Proof. (i) Take $x \in B_r(e_0) \cap B_s(e_1)$. By triangle inequality, $d(e_0, e_1) \leq d(e_0, x) + d(x, e_1) \leq r + s$.

(ii) Let $\delta = d(e_0, e_1)$. If $r = 0$ or $s = 0$ the statement is trivial. Otherwise define $x = e_0 + \frac{r}{r+s}(e_1 - e_0)$. Then

$$\|x - e_0\| = \frac{r}{r+s}\delta \leq r, \qquad \|x - e_1\| = \frac{s}{r+s}\delta \leq s,$$

so x lies in the intersection. □

Interpretation 3.7. Two epistemic states can be reconciled—there exists a membership function compatible with both tolerances—if and only if the distance between their centers does not exceed the sum of their tolerances. This quantifies when expert opinions are mutually consistent.

Corollary 3.2 (Minimal Common Tolerance). *In a normed linear space, the smallest $r \geq 0$ such that $B_r(e_0) \cap B_r(e_1) \neq \emptyset$ is $r^* = \frac{1}{2}d(e_0, e_1)$.*

Proof. From the intersection condition with $r = s$, intersection occurs if and only if $d(e_0, e_1) \leq 2r$, i.e., $r \geq \frac{1}{2}d(e_0, e_1)$. □

Interpretation 3.8. If two experts have different reference assessments, the minimum common tolerance needed for their balls to overlap is exactly half the distance between their references. This provides a quantitative measure of how much they must broaden their uncertainty to reach agreement.

3.5 Equivalence to Uniform-Radius Interval-Valued Fuzzy Sets

Ball fuzzy sets are closely related to interval-valued fuzzy sets (IVFS), but the relationship is not one of equality. A ball fuzzy set induces an IVFS via its canonical envelope. However, not every IVFS arises from a ball.

Definition 3.3 (Interval-Valued Fuzzy Set)**.** An *interval-valued fuzzy set* on X is a pair of functions (A_-, A_+) with $A_-(x) \le A_+(x)$ for all $x \in X$.

Theorem 3.4 (Characterization of IVFS Arising from Balls)**.** *An IVFS (A_-, A_+) arises from a ball fuzzy set if and only if there exist $\mu_0 \in [0,1]^X$ and $r \ge 0$ such that*

$$A_-(x) = \max(0, \mu_0(x) - r), \qquad A_+(x) = \min(1, \mu_0(x) + r) \quad \forall x \in X.$$

Equivalently, the lengths $A_+(x) - A_-(x)$ are constant (equal to $2r$) except where truncated by the bounds 0 or 1, and the midpoints are independent of x when not truncated.

Proof. If the IVFS arises from a ball, the condition is immediate. Conversely, if the lengths are constant (ignoring truncation) and the midpoints are independent of x, set $\mu_0(x)$ equal to the common midpoint; the truncation at boundaries automatically yields the given A_-, A_+. □

Corollary 3.3. *Ball fuzzy sets are a* proper subset *of interval-valued fuzzy sets. There exist IVFS that are not ball fuzzy sets.*

Example 3.1. Let $X = \{x_1, x_2\}$ and define $A_-(x_1) = 0.2$, $A_+(x_1) = 0.8$, $A_-(x_2) = 0.3$, $A_+(x_2) = 0.7$. The lengths are 0.6 and 0.4, which are not equal. No ball fuzzy set can produce this IVFS.

Interpretation 3.9. The uniform-radius constraint is the defining characteristic of ball fuzzy sets. It reflects the global nature of the epistemic tolerance: uncertainty is modeled by a single parameter r that applies uniformly to all elements, rather than allowing different margins at different coordinates. This global constraint is precisely what enables the algebraic operations developed in Section 4 to preserve the ball structure.

3.6 Coordinate-Wise Characterization

For finite universes, the Chebyshev ball has a simple coordinate-wise characterization that is essential for computation.

Theorem 3.5 (Coordinate-Wise Characterization)**.** *For any $\mu_0 \in [0,1]^X$ and $r \ge 0$,*

$$B_r(\mu_0) = \left\{ \mu \in [0,1]^X \;\middle|\; \mu(x) \in [\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)] \; \forall x \in X \right\}.$$

Equivalently,

$$B_r(\mu_0) = \prod_{x \in X} [\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)].$$

Proof. The condition $\|\mu - \mu_0\|_\infty \le r$ is equivalent to $|\mu(x) - \mu_0(x)| \le r$ for every x. Intersecting with the requirement $\mu(x) \in [0,1]$ yields the interval bounds. The product representation follows because the constraints are independent across coordinates. □

Remark 3.2 (Finite Universe Assumption)**.** Throughout this work we assume X is finite. This is standard in practical fuzzy set applications and simplifies the mathematical exposition. For infinite universes, $\max$ must be replaced by $\sup$ in the definition of the Chebyshev metric. The results extend naturally with appropriate technical adjustments.

Summary

This section has established the foundational structural properties of ball fuzzy sets. We have:

1. Defined ball fuzzy sets as the geometric realization of the PBF in $[0,1]^X$ equipped with the Chebyshev metric.
2. Derived the canonical interval-valued representation and the three-way decomposition into belief, epistemic uncertainty, and disbelief.
3. Established the fundamental axioms (Reference and Bound, Closure) and derived monotonicity and degeneracy.
4. Proved key geometric properties: convexity, translation invariance, Lipschitz stability, diameter bound, compactness, center identifiability, and intersection conditions.
5. Characterized ball fuzzy sets as a proper subset of interval-valued fuzzy sets—specifically, uniform-radius IVFS.
6. Provided the coordinate-wise characterization for finite universes.

These properties provide the geometric and epistemic foundation for the algebraic operations developed in Section 4.

4. Algebraic Structure of Ball Fuzzy Sets

Section 3 established the fundamental geometric properties of ball fuzzy sets: their definition as Chebyshev balls in $[0,1]^X$, their canonical interval-valued representation, the three-way decomposition into belief, epistemic uncertainty, and disbelief, and their characterization as uniform-radius interval-valued fuzzy sets. This section develops the natural algebraic operations on ball fuzzy sets: complement, intersection, and union.

A ball fuzzy set is, fundamentally, a *set* of membership functions—the set of all functions within tolerance r of the reference μ_0. Consequently, the most natural definitions of algebraic operations are set-theoretic:

$$\begin{aligned} \mathcal{B}^c &= \{\mu^c \mid \mu \in \mathcal{B}\}, \\ \mathcal{B}_1 \cap \mathcal{B}_2 &= \{\mu \mid \mu \in \mathcal{B}_1 \text{ and } \mu \in \mathcal{B}_2\}, \\ \mathcal{B}_1 \cup \mathcal{B}_2 &= \{\mu \mid \mu \in \mathcal{B}_1 \text{ or } \mu \in \mathcal{B}_2\}, \end{aligned}$$

where $\mu^c(x) = 1 - \mu(x)$ is the standard fuzzy complement.

However, the result of such an operation is not guaranteed to be a ball fuzzy set. A central theme of this section is to characterize exactly when the result is again a ball, and when it is not, to identify the "best" ball that approximates it—typically the smallest ball containing the set (for union) or the largest ball contained in it (for intersection). All operations are derived from the underlying geometry, not postulated arbitrarily, and each comes with a clear epistemic interpretation.

4.1 Complement Operation

The complement of a fuzzy set negates the concept: membership in the complement corresponds to non-membership in the original. For ball fuzzy sets, this operation has a particularly elegant property.

Definition 4.1 (Complement)**.** Let $\mathcal{B} = B_r(\mu_0)$ be a ball fuzzy set. Its *complement* is

$$\mathcal{B}^c = \{\mu^c \mid \mu \in \mathcal{B}\},$$

where $\mu^c(x) = 1 - \mu(x)$ for all $x \in X$.

Theorem 4.1 (Complement Preserves Ball Structure)**.** *For any* $\mu_0 \in [0,1]^X$ *and* $r \geq 0$,

$$\big(B_r(\mu_0)\big)^c = B_r(\mu_0^c),$$

where $\mu_0^c(x) = 1 - \mu_0(x)$.

Proof. For any $\mu \in B_r(\mu_0)$, we have $\|\mu - \mu_0\|_\infty \le r$. Since

$$|\mu^c(x) - \mu_0^c(x)| = |1 - \mu(x) - (1 - \mu_0(x))| = |\mu(x) - \mu_0(x)|$$

for every x, it follows that $\|\mu^c - \mu_0^c\|_\infty = \|\mu - \mu_0\|_\infty \le r$. Thus $\mu^c \in B_r(\mu_0^c)$. The converse is symmetric. □

Corollary 4.1 (Involution). $(\mathcal{B}^c)^c = \mathcal{B}$ *for every ball fuzzy set* $\mathcal{B}$.

Interpretation 4.1. Complement corresponds to negating the concept while preserving the same tolerance. If μ_0 is the reference assessment for a fuzzy set A, then μ_0^c is the natural reference for the complement of A, and the admissible region around it has exactly the same radius. This symmetry is a direct consequence of the Chebyshev metric's invariance under the transformation $\mu \mapsto 1 - \mu$.

4.2 Intersection Operation

The intersection of two fuzzy sets corresponds to the conjunction of concepts: an element belongs to the intersection to the degree that it belongs to both. For ball fuzzy sets, the set-theoretic intersection of two balls is not always a ball. We must therefore distinguish between the *set-theoretic intersection* and the *intersection ball*.

Definition 4.2 (Set-Theoretic Intersection). Let $\mathcal{B}_1 = B_r(\mu_0)$ and $\mathcal{B}_2 = B_s(\nu_0)$ be two ball fuzzy sets. Their *set-theoretic intersection* is

$$\mathcal{B}_1 \cap \mathcal{B}_2 = \{\mu \mid \mu \in \mathcal{B}_1 \text{ and } \mu \in \mathcal{B}_2\}.$$

Recall that each ball is a Cartesian product of intervals:

$$B_r(\mu_0) = \prod_{x \in X} I_x, \qquad B_s(\nu_0) = \prod_{x \in X} J_x,$$

where

$$I_x = [\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)], \qquad J_x = [\max(0, \nu_0(x) - s),\ \min(1, \nu_0(x) + s)].$$

The intersection is therefore the Cartesian product of the per-coordinate intersections:

$$\mathcal{B}_1 \cap \mathcal{B}_2 = \prod_{x \in X} (I_x \cap J_x).$$

For this product to be a Chebyshev ball, all intervals $I_x \cap J_x$ must have the same length (up to truncation at the boundaries 0 and 1), and their midpoints must be independent of x. The following theorem characterizes exactly when this occurs.

Theorem 4.2 (Intersection Criterion). *Let* $\mathcal{B}_1 = B_r(\mu_0)$ *and* $\mathcal{B}_2 = B_s(\nu_0)$ *with* $\mathcal{B}_1 \cap \mathcal{B}_2 \neq \emptyset$. *Then* $\mathcal{B}_1 \cap \mathcal{B}_2$ *is itself a ball fuzzy set if and only if*

$$d_\infty(\mu_0, \nu_0) \le |r - s|.$$

When this condition holds, we have

$$B_r(\mu_0) \cap B_s(\nu_0) = B_{\min(r,s)}(\mu^*),$$

where μ^* *is the unique center of the resulting ball.*

Proof. The intersection interval at coordinate x is

$$I_x \cap J_x = \big[\max(\mu_0(x) - r, \nu_0(x) - s, 0),\ \min(\mu_0(x) + r, \nu_0(x) + s, 1)\big].$$

Its length is

$$\ell_x = \min(\mu_0(x) + r, \nu_0(x) + s, 1) - \max(\mu_0(x) - r, \nu_0(x) - s, 0).$$

For the product $\prod_x (I_x \cap J_x)$ to be a Chebyshev ball, we require all ℓ_x to be equal and the midpoints to be independent of x. A straightforward calculation shows that ℓ_x is independent of x iff

$$\max_x |\mu_0(x) - \nu_0(x)| \le |r - s|,$$

i.e., $d_\infty(\mu_0, \nu_0) \le |r - s|$. When this holds, the common length is $2\min(r, s)$ (after accounting for boundary truncation), and the midpoints become constant, yielding the ball $B_{\min(r,s)}(\mu^*)$. $\square$

Corollary 4.2. *If $r = s$, then $\mathcal{B}_1 \cap \mathcal{B}_2$ is a ball if and only if $\mu_0 = \nu_0$. Otherwise, the intersection is a hyperrectangle with varying side lengths.*

Interpretation 4.2. Two epistemic states are compatible enough that their intersection is again a ball exactly when their centers are close relative to the difference of their tolerances. In that case, the common admissible region is described by the smaller tolerance around a suitably chosen center. If the centers are too far apart, the intersection becomes a hyperrectangle with possibly varying side lengths and cannot be captured by a single uniform tolerance.

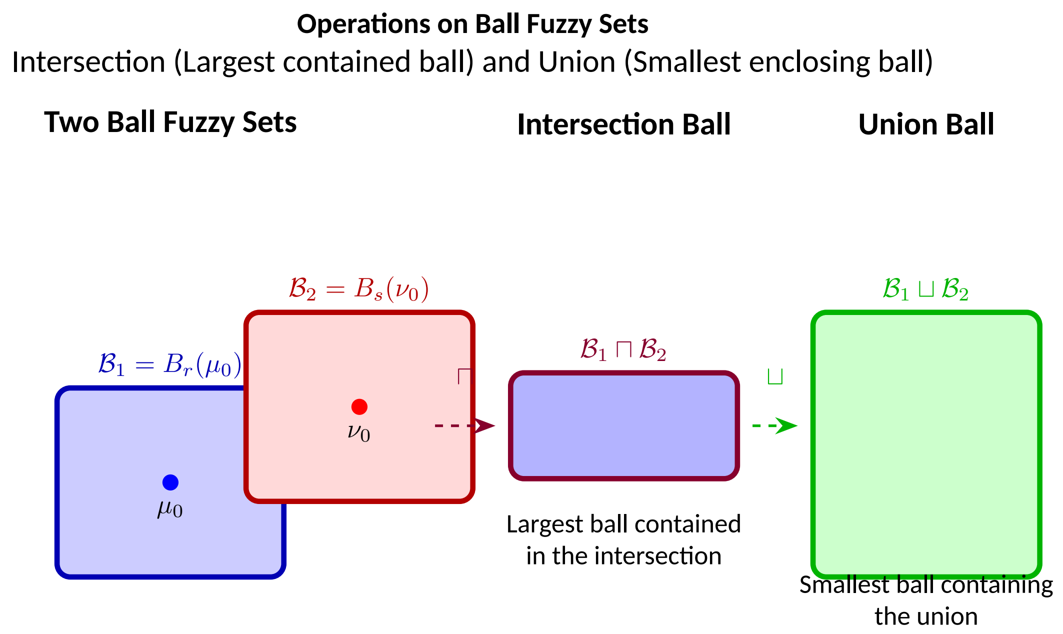

4.2.1 *The Intersection Ball*

When the set-theoretic intersection is not a ball, we may still wish to represent the common admissible region by a ball fuzzy set that is *contained* in the intersection. The natural choice is the largest such ball.

Definition 4.3 (Intersection Ball). For two ball fuzzy sets $\mathcal{B}_1, \mathcal{B}_2$ with non-empty intersection, their *intersection ball* $\mathcal{B}_1 \sqcap \mathcal{B}_2$ is the largest ball (with respect to inclusion) that is contained in $\mathcal{B}_1 \cap \mathcal{B}_2$.

Theorem 4.3 (Existence and Form of the Intersection Ball). *For $\mathcal{B}_1 = B_r(\mu_0)$ and $\mathcal{B}_2 = B_s(\nu_0)$ with non-empty intersection, the intersection ball exists and is given by $B_\rho(\xi)$, where*

$$\rho = \frac{1}{2} \min_{x \in X} \big[\min(\mu_0(x) + r, \nu_0(x) + s, 1) - \max(\mu_0(x) - r, \nu_0(x) - s, 0) \big],$$

and ξ is defined componentwise by

$$\xi(x) = \frac{\max(\mu_0(x) - r, \nu_0(x) - s, 0) + \min(\mu_0(x) + r, \nu_0(x) + s, 1)}{2}.$$

Proof. The intersection $\mathcal{B}_1 \cap \mathcal{B}_2$ is the Cartesian product of the coordinate projections. Any ball contained in this product must satisfy $\mathcal{B}(x) \subseteq \mathcal{B}_1(x) \cap \mathcal{B}_2(x)$ for every x. Hence its radius is at most half the minimum length of these intervals. The center ξ is the Chebyshev center of the hyperrectangle $\prod_x (I_x \cap J_x)$; it maximizes the minimal distance to the boundaries, yielding the largest possible radius. $\square$

Interpretation 4.3. When two epistemic states are not perfectly compatible, the largest region that can be described by a uniform tolerance and is guaranteed to be admissible under both is given by the intersection ball. Its radius is limited by the coordinate where the admissible intervals are narrowest—the "weakest link" in the consensus. This provides a conservative but principled way to combine conflicting information while maintaining the uniform-tolerance structure.

4.3 Union Operation

The union of two fuzzy sets corresponds to the disjunction of concepts: an element belongs to the union to the degree that it belongs to either. For ball fuzzy sets, the set-theoretic union is almost never a ball. We therefore define the union operation as the smallest ball that contains both given balls.

Definition 4.4 (Union Ball)**.** For two ball fuzzy sets $\mathcal{B}_1, \mathcal{B}_2$, their *union ball* $\mathcal{B}_1 \sqcup \mathcal{B}_2$ is the smallest ball (with respect to inclusion) that contains $\mathcal{B}_1 \cup \mathcal{B}_2$.

Theorem 4.4 (Formula for the Union Ball)**.** *Let $\mathcal{B}_1 = B_r(\mu_0)$ and $\mathcal{B}_2 = B_s(\nu_0)$ be two ball fuzzy sets on a finite universe X. For each $x \in X$, define*

$$a(x) = \min\big(\max(0, \mu_0(x) - r),\ \max(0, \nu_0(x) - s)\big),$$
$$b(x) = \max\big(\min(1, \mu_0(x) + r),\ \min(1, \nu_0(x) + s)\big).$$

Then the union ball $\mathcal{B}_1 \sqcup \mathcal{B}_2$ is $B_R(\mu_U)$, where

$$\mu_U(x) = \frac{a(x) + b(x)}{2}, \qquad R = \frac{1}{2}\max_{x \in X}\big(b(x) - a(x)\big).$$

In other words, the union ball is the Chebyshev ball whose interval at each x is exactly the convex hull of the two original intervals.

Proof. The union $\mathcal{B}_1 \cup \mathcal{B}_2$ is contained in the product $\prod_x [a(x), b(x)]$, and this product is the smallest axis-aligned hyperrectangle containing the union. The minimal enclosing ball of a hyperrectangle in the Chebyshev metric is the ball centered at the componentwise midpoint with radius half the maximum side length. Thus $B_R(\mu_U)$ as defined is the smallest Chebyshev ball containing $\mathcal{B}_1 \cup \mathcal{B}_2$. $\square$

Interpretation 4.4. The union of two epistemic states combines all assessments admissible under either source. The resulting tolerance must be large enough to cover both original ranges; the new center is the midpoint of the extremes. This operation inevitably introduces new membership functions that were not admissible under either original state—a necessary cost if we insist on representing the combined knowledge by a single uniform tolerance. In practice, this corresponds to taking a more conservative stance when aggregating disparate sources of information.

4.4 Lattice Properties

The operations $\sqcap$ (intersection ball), $\sqcup$ (union ball), and complement satisfy the following algebraic laws.

Theorem 4.5 (Lattice Properties)**.** *For any ball fuzzy sets $\mathcal{A}, \mathcal{B}, \mathcal{C}$:*

(i) Commutativity: $\mathcal{A} \sqcap \mathcal{B} = \mathcal{B} \sqcap \mathcal{A}$, $\mathcal{A} \sqcup \mathcal{B} = \mathcal{B} \sqcup \mathcal{A}$.

(ii) Idempotence: $\mathcal{A} \sqcap \mathcal{A} = \mathcal{A}$, $\mathcal{A} \sqcup \mathcal{A} = \mathcal{A}$.

(iii) Absorption: $\mathcal{A} \sqcap (\mathcal{A} \sqcup \mathcal{B}) = \mathcal{A}$, $\mathcal{A} \sqcup (\mathcal{A} \sqcap \mathcal{B}) = \mathcal{A}$.

(iv) Associativity: $(\mathcal{A} \sqcap \mathcal{B}) \sqcap \mathcal{C} = \mathcal{A} \sqcap (\mathcal{B} \sqcap \mathcal{C})$ and $(\mathcal{A} \sqcup \mathcal{B}) \sqcup \mathcal{C} = \mathcal{A} \sqcup (\mathcal{B} \sqcup \mathcal{C})$ hold whenever all intermediate intersections/unions are themselves balls.

(v) Partial distributivity: The distributive laws hold when all involved balls lie in the interior $(0,1)^X$ and all intermediate intersections are balls. Boundary truncation can cause failure, reflecting the behaviour of interval arithmetic.

Sketch. Commutativity follows from symmetry of the definitions. Idempotence and absorption are immediate from the definitions.

Associativity for $\sqcap$: If $(\mathcal{A} \cap \mathcal{B}) \cap \mathcal{C}$ is a ball, then $\mathcal{A} \sqcap (\mathcal{B} \sqcap \mathcal{C})$ is the largest ball contained in $\mathcal{A} \cap (\mathcal{B} \cap \mathcal{C})$, which is the same ball. Similarly for union. However, if intermediate intersections are not balls, associativity may fail.

Distributivity holds under the stated interior and intersection conditions because the operations reduce to set-theoretic intersection and union on intervals of constant length. Boundary truncation at 0 and 1 can break distributivity, as in interval arithmetic. □

Corollary 4.3. *The family of ball fuzzy sets, together with the operations $\sqcap$ and $\sqcup$, forms a lattice with a well-defined complement satisfying involution. It is not a Boolean algebra because the law of excluded middle fails: $\mathcal{A} \sqcup \mathcal{A}^c$ is generally not the whole space $[0, 1]^X$.*

Interpretation 4.5. These algebraic properties ensure that ball fuzzy sets behave in a predictable and coherent way when combined. Commutativity and associativity guarantee that the order of combination does not matter. Absorption reflects the natural hierarchy between a set and its superset/subset. The failure of full distributivity at boundaries is consistent with interval arithmetic and does not hinder practical application.

4.5 Behavior at Different Center Types

The Principle of Bounded Fuzziness allows the reference assessment μ_0 to be any function from X to $[0, 1]$. However, three particular classes of centers are natural and significant:

- *Crisp centers:* $\mu_0 : X \to \{0, 1\}$, corresponding to characteristic functions of classical sets.
- *Ternary centers:* $\mu_0 : X \to \{0, 0.5, 1\}$, introducing a neutral value 0.5 representing complete indeterminacy.
- *Fully fuzzy centers:* $\mu_0 : X \to [0, 1]$, the fully graded case.

These classes are nested:

$$\{\text{crisp}\} \subset \{\text{ternary}\} \subset \{\text{fuzzy}\}.$$

Each represents a different level of granularity in the reference assessment, and the resulting ball fuzzy sets inherit distinctive properties from their centers.

4.5.1 Crisp Centers

When $\mu_0 = \chi_A$ for some $A \subseteq X$, the ball $B_r(\chi_A)$ takes a particularly simple form:

$$\underline{\mu}(x) = \max(0, \chi_A(x) - r), \qquad \overline{\mu}(x) = \min(1, \chi_A(x) + r).$$

Proposition 4.1 (Crisp Centers Generate Symmetric Intervals). *For a crisp center χ_A, for each $x \in X$:*

- *If $x \in A$: the admissible interval is $[1 - r, 1]$ (truncated at 1);*
- *If $x \notin A$: the admissible interval is $[0, r]$ (truncated at 0).*

The intervals are symmetric about the center value (1 for members, 0 for non-members).

Interpretation 4.6. When the reference assessment is crisp, uncertainty is modeled by allowing membership grades to deviate from the crisp values 0 or 1 by at most r. This yields a simple "band" around the classical set: elements originally in A can have membership as low as $1 - r$, while elements originally not in A can have membership as high as r.

4.5.2 *Ternary Centers*

Ternary centers introduce the value 0.5 as a third possibility, representing complete epistemic indifference—neither membership nor non-membership is favored.

Proposition 4.2 (Ternary Centers Create Three Zone Types). *For a ternary center μ_0, define*

$$A_0 = \{x \mid \mu_0(x) = 0\}, \qquad A_{0.5} = \{x \mid \mu_0(x) = 0.5\}, \qquad A_1 = \{x \mid \mu_0(x) = 1\}.$$

Then for any $r \geq 0$, the admissible intervals are:

- $x \in A_0$: $[0, \min(1, r)]$;
- $x \in A_{0.5}$: $[\max(0, 0.5 - r), \min(1, 0.5 + r)]$;
- $x \in A_1$: $[\max(0, 1 - r), 1]$.

Interpretation 4.7. The ternary center introduces a third qualitative state: elements judged as completely indeterminate (0.5) generate intervals centered at the midpoint of the unit interval. For small r (e.g., $r < 0.5$), these intervals are strictly contained in $(0, 1)$, never reaching the extremes. This captures the idea that genuine uncertainty about membership should not be resolved by default to either 0 or 1.

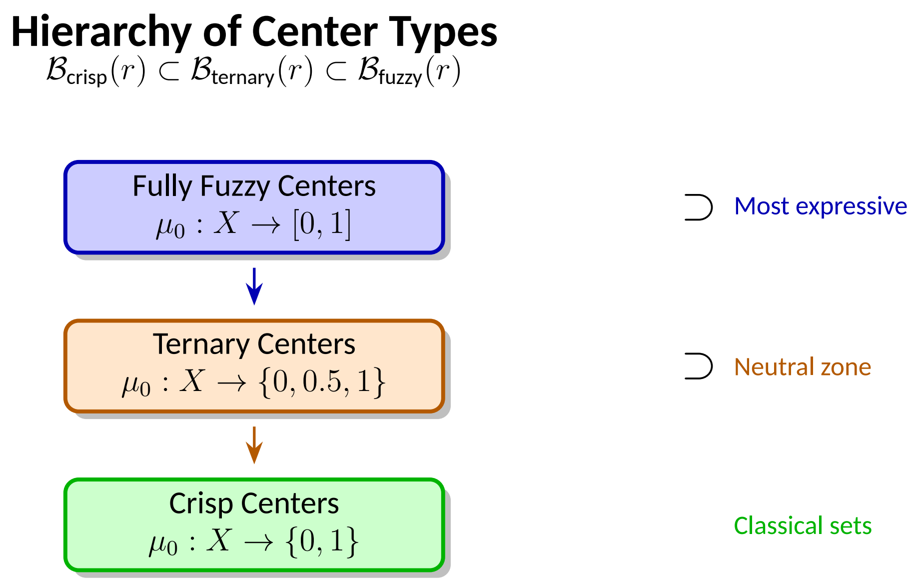

Complement preserves center type; intersection and union may generate higher types.

4.5.3 *Fully Fuzzy Centers*

When μ_0 can take any value in $[0, 1]$, the admissible intervals are simply

$$[\max(0, \mu_0(x) - r),\ \min(1, \mu_0(x) + r)].$$

No special symmetry or structure is imposed beyond the uniform tolerance. This is the most general case.

4.5.4 *Nesting of Center Classes*

The three classes form a natural hierarchy in terms of expressive power and structural constraints.

Theorem 4.6 (Nesting of Classes). *Let $\mathcal{B}_{crisp}(r)$, $\mathcal{B}_{ternary}(r)$, and $\mathcal{B}_{fuzzy}(r)$ denote the families of ball fuzzy sets with radius r and centers of the respective types. Then*

$$\mathcal{B}_{crisp}(r) \subset \mathcal{B}_{ternary}(r) \subset \mathcal{B}_{fuzzy}(r)$$

for any fixed $r > 0$. The inclusions are strict.

Sketch. Any crisp center is also ternary (since $\{0,1\} \subset \{0,0.5,1\}$), and any ternary center is fuzzy, giving the inclusions. For strictness: a ball with a ternary center having $A_{0.5} \neq \emptyset$ and $r < 0.5$ produces intervals strictly contained in $(0,1)$ for elements in $A_{0.5}$. Such intervals cannot arise from a crisp center, where intervals always touch either 0 or 1. A fully fuzzy center with a value not in $\{0,0.5,1\}$, e.g., $\mu_0(x) = 0.3$, yields an interval with midpoint 0.3, which no ternary center can produce. □

4.5.5 Algebraic Closure Properties

Do the operations preserve the type of the center? The answer reveals further structure.

Theorem 4.7 (Complement Preserves Center Type)**.** *If $\mathcal{B}$ has a crisp (resp. ternary, fuzzy) center, then its complement $\mathcal{B}^c$ also has a crisp (resp. ternary, fuzzy) center.*

Proof. The complement center is μ_0^c. If μ_0 takes values in $\{0,1\}$, so does μ_0^c; similarly for $\{0,0.5,1\}$ and $[0,1]$. □

Theorem 4.8 (Intersection and Union Do Not Preserve Center Type)**.** *The intersection ball and union ball of two crisp-centered balls need not have crisp centers; they may be ternary or even fully fuzzy. Similarly, the intersection of two ternary-centered balls may yield a fuzzy-centered ball.*

Example 4.1. Take $X = \{x\}$, $\mathcal{B}_1 = B_r(0)$ (crisp center at 0), $\mathcal{B}_2 = B_r(1)$ (crisp center at 1). Their union ball has center 0.5 and radius $r + 0.5$ (assuming r small enough). The center 0.5 is ternary but not crisp.

Interpretation 4.8. Combining epistemic states from different crisp sources naturally generates intermediate reference values. This reflects the idea that reconciling opposing viewpoints yields a compromise that may itself be indeterminate (ternary) or fully graded.

Summary

This section has developed the algebraic structure of ball fuzzy sets. We have:

1. Defined the complement operation and proved it preserves the ball structure.
2. Characterized when the set-theoretic intersection of two balls is itself a ball, and defined the intersection ball as the largest ball contained in the intersection when it is not.
3. Defined the union ball as the smallest ball containing the set-theoretic union.
4. Established the lattice properties of the resulting algebra.
5. Analyzed the behavior of three natural center types—crisp, ternary, and fully fuzzy—establishing their nesting and algebraic closure properties.

These algebraic operations provide a principled way to combine and manipulate ball fuzzy sets while preserving the geometric structure that gives the framework its coherence. The following section demonstrates the framework in a realistic application.

Remark 4.1 (Comparison to Interval-Valued Fuzzy Set Operations)**.** The operations on ball fuzzy sets refine the usual IVFS operations by imposing the additional uniformity constraint. In particular, the interval-wise intersection of two ball fuzzy sets may not have uniform length; the intersection ball adjusts the result to the largest uniform-radius set contained in it. This reflects the principle that combining epistemic states should yield a representation of the same structured type whenever possible.

5. Empirical Validation of the Principle of Bounded Fuzziness

The Principle of Bounded Fuzziness (PBF) asserts that when multiple assessors possess shared knowledge about an object, their assessments do not disperse arbitrarily throughout the uncertainty space. Instead, the induced uncertainty remains concentrated within a bounded region around a collective reference.

The objective of this section is therefore not to verify a collection of independent empirical observations. Rather, it is to determine whether bounded epistemic disagreement is an observable phenomenon across different uncertainty-generation mechanisms.

To this end, we employ two datasets representing distinct manifestations of uncertainty:

1. CIFAR-10H [15], which captures human uncertainty through empirical probability distributions over image classes.
2. CURVAS [16], which represents expert disagreement in medical image assessment.

The experiments are organized around four fundamental questions implied by the PBF:

1. Does a probabilistically bounded disagreement region exist?
2. Does shared knowledge constrain disagreement relative to random knowledge?
3. Does the estimated bound remain stable under increasing sample sizes and resampling?
4. Is the phenomenon observable across heterogeneous domains?

Table 1
Datasets used for empirical validation.

Dataset	Domain	Status	Samples	Assessors
CIFAR-10H	Image Classification	Real	500	Human Annotators
CURVAS	Medical Imaging	Simulation	90	Expert Raters

5.1 *Existence of a Probabilistically Bounded Region*

The primary prediction of the PBF is the existence of a finite disagreement radius surrounding the reference assessment.

For a desired confidence level α, the corresponding coverage radius is defined by

$$r_\alpha = \inf\{r : P(D \le r) \ge \alpha\},$$

where $D = |\mu_j(x) - \mu_0(x)|$ denotes the disagreement between an assessment and the reference value.

Table 2 reports the estimated coverage radii.

Table 2
Coverage radii estimated from empirical disagreement distributions.

Dataset	r_{90}	r_{95}	r_{99}
CIFAR-10H	0.1600	0.2521	0.8491
CURVAS	0.1581	0.2062	0.3092

Several observations are noteworthy.

First, finite radii exist at all confidence levels for both datasets. Second, the estimated radii are accompanied by relatively narrow bootstrap confidence intervals, indicating statistical stability. Third, the majority of disagreements are concentrated within remarkably small neighborhoods of the reference assessment.

These findings provide direct empirical support for the central prediction of the PBF: epistemic disagreement is not unbounded but instead occupies a finite probabilistic region.

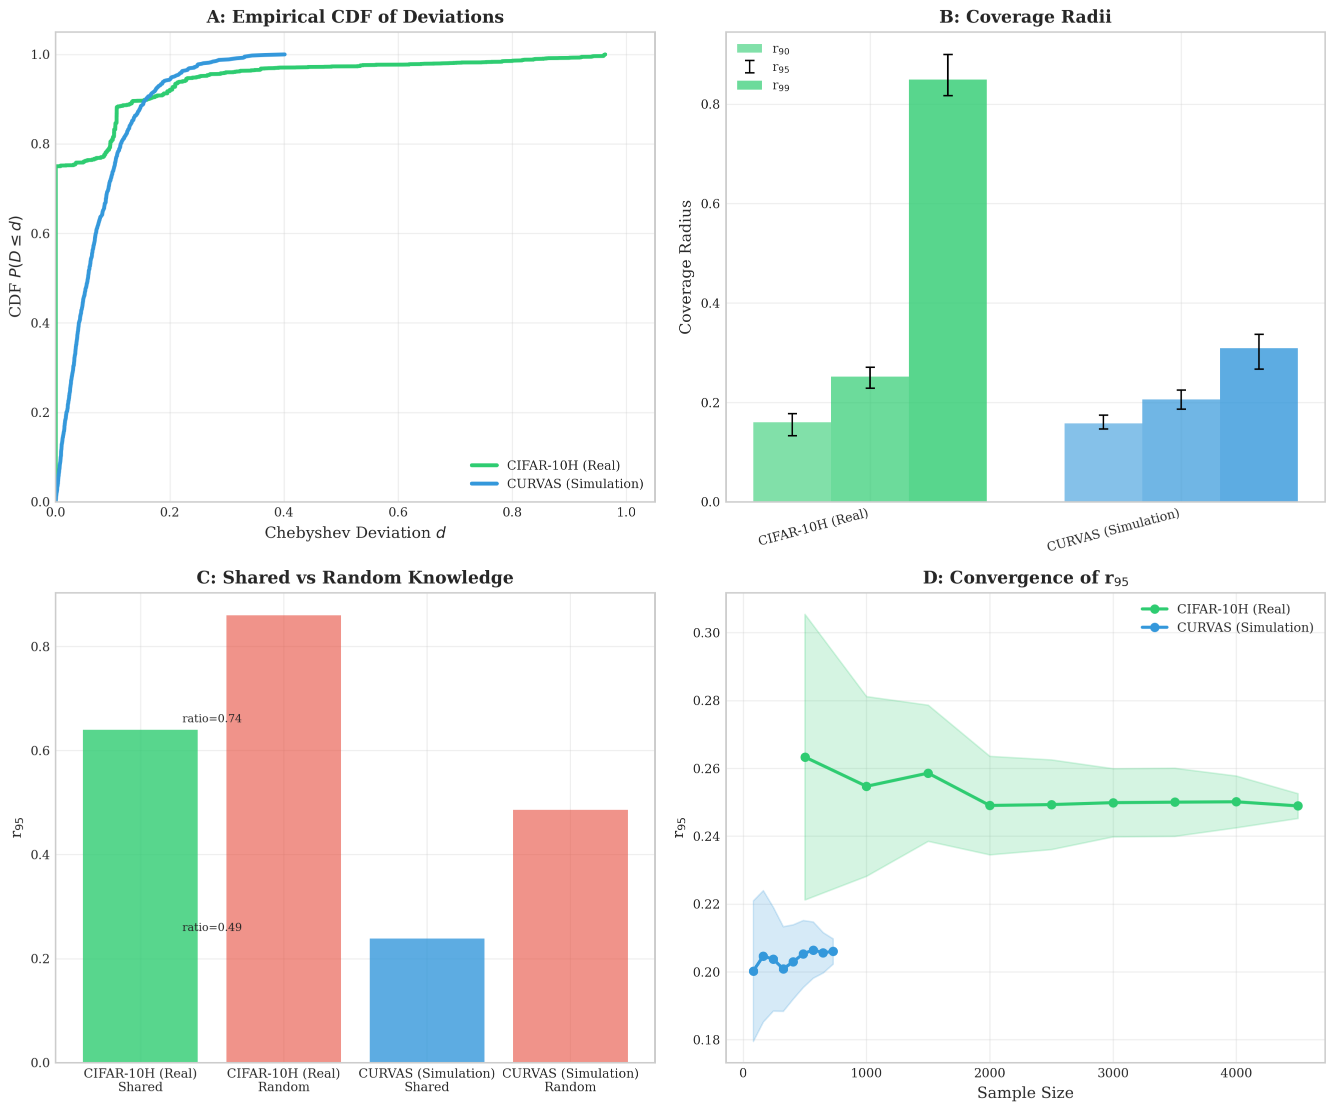

Fig. 1. Empirical validation of the Principle of Bounded Fuzziness — Properties 1–4. **A**: Empirical CDF of Chebyshev deviations $D = |\mu_j(x) - \mu_0(x)|$. The CDF (cumulative distribution function) shows the probability $P(D \leq d)$ that a deviation does not exceed a given value d. A steep rise near zero indicates that the vast majority of assessments are tightly clustered around the reference. The horizontal dashed line at $y = 0.95$ intersects the curves at $d = r_{95}$, the radius capturing 95% of all deviations. **B**: Coverage radii r_{90}, r_{95}, r_{99} with bootstrap 95% confidence intervals (error bars). **C**: Shared-knowledge disagreement is significantly smaller than random-knowledge disagreement (ratios: CIFAR-10H = 0.744, CURVAS = 0.489; $p < 0.001$). **D:** r_{95} converges as sample size increases ($CV < 0.05$).

Figure 1 visualizes the first four properties. Panel A displays the empirical cumulative distribution function (CDF) of Chebyshev deviations. The CDF $P(D \leq d)$ gives the proportion of deviations not exceeding a threshold d. The steep rise near zero confirms that the majority of assessments are tightly concentrated around the reference. The horizontal line at $y = 0.95$ intersects the curves at $d = r_{95}$, providing a direct visual interpretation of the coverage radii reported in Table 2: 95% of all deviations fall within $r_{95} = 0.252$ for CIFAR-10H and $r_{95} = 0.206$ for CURVAS. Panel B displays the coverage radii r_{90}, r_{95}, r_{99} with their 95% bootstrap confidence intervals. Panel C demonstrates that shared knowledge produces significantly smaller disagreement than random knowledge, with ratios of 0.744 for CIFAR-10H and 0.489 for CURVAS ($p < 0.001$). Panel D reveals the convergence of r_{95} as the sample size increases, with coefficients of variation well below 0.05.

5.2 *Knowledge as a Constraint on Disagreement*

The philosophical foundation of the PBF is that knowledge acts as a constraint on disagreement.

To evaluate this claim, the observed disagreement distributions were compared against null models in which the knowledge structure was deliberately destroyed through randomization.

If the PBF is valid, disagreement generated under shared knowledge should be significantly smaller than disagreement generated under random knowledge.

The results are summarized in Table 3.

Table 3
Comparison between shared-knowledge and random-knowledge disagreement.

Dataset	Shared r_{95}	Random r_{95}	Ratio
CIFAR-10H	0.6400	0.8600	0.744
CURVAS	0.2385	0.4875	0.489

For both datasets the reduction in disagreement is highly significant ($p < 0.001$).

The effect is particularly pronounced for the medical dataset, where the disagreement radius under shared knowledge is less than half of that observed under randomization.

This result constitutes the strongest empirical evidence supporting the PBF. It directly confirms the principle's central claim that shared knowledge constrains the extent of epistemic disagreement.

5.3 *Stability and Convergence of the Bound*

A fundamental requirement of any uncertainty principle is that its predictions should not depend critically on the particular sample used for estimation. If the estimated disagreement bound varies substantially when different subsets of the data are considered, or when the sample is resampled, then the observed boundedness may be an artifact of sampling rather than a genuine structural property of the underlying uncertainty. Conversely, if the estimated bound remains stable across different samples and converges as more data becomes available, this provides strong evidence that boundedness is a real phenomenon.

To investigate this issue, we conducted two complementary analyses: (i) a convergence analysis examining how the estimated radius r_{95} changes as the number of assessors increases, and (ii) a bootstrap stability analysis examining how r_{95} varies across resampled datasets.

Convergence Analysis. The convergence analysis addresses the question: as we incorporate assessments from an increasing number of experts, does the estimated disagreement bound stabilize, or does it continue to grow without limit?

The procedure is as follows. For each dataset, we randomly sampled subsets of the available observations at increasing sizes. Specifically, we used sample sizes ranging from 10% to 90% of the total number of elements in increments of 10%. For each subset size, we performed 30 bootstrap replications: we randomly selected that many elements without replacement, computed the Chebyshev deviations from the leave-one-out consensus, and estimated r_{95} as the 95th percentile of the resulting deviation distribution. We then computed the mean and standard deviation of r_{95} across the 30 replications for each subset size.

This procedure yields a sequence of estimated radii $r_{95}(m)$ as a function of the sample size m. If the PBF correctly captures a genuine structural property, this sequence should converge to a finite value r^* as m increases. Mathematically, we expect

$$\lim_{m \to \infty} r_{95}(m) = r^* < \infty.$$

The rate of convergence provides additional insight: rapid convergence indicates that the bound is determined by the underlying structure rather than by sample-specific fluctuations. We quantify convergence using the coefficient of variation (CV) computed over the final three sample sizes:

$$\text{CV} = \frac{\text{std}(r_{95}^{(k-2)}, r_{95}^{(k-1)}, r_{95}^{(k)})}{\text{mean}(r_{95}^{(k-2)}, r_{95}^{(k-1)}, r_{95}^{(k)})}.$$

A CV below 0.05 (i.e., 5%) is interpreted as evidence of convergence.

For CIFAR-10H, the sequence $r_{95}(m)$ converged rapidly, stabilizing at $r_{95} = 0.2495$ with a coefficient of variation of 0.0017 over the final three sample sizes. For CURVAS, the sequence converged to $r_{95} = 0.2058$ with a coefficient of variation of 0.0031. These results are summarised in Table 4.

Table 4
Convergence of r_{95} with increasing sample size.

Dataset	Final r_{95}	CV (last 3)	Converged
CIFAR-10H	0.2495	0.0017	Yes
CURVAS	0.2058	0.0031	Yes

These findings demonstrate that the estimated bound is not an artifact of the particular sample used. As more data becomes available, the radius stabilizes rapidly, indicating that the underlying disagreement structure is consistent and bounded.

Bootstrap Stability Analysis. While the convergence analysis examines stability as the sample size increases, the bootstrap stability analysis examines stability under resampling of the existing data. This addresses a different question: if we were to repeat the experiment with a different random sample of the same size, would we obtain a substantially different estimate of r_{95}?

The procedure is as follows. We generated 500 bootstrap samples by sampling with replacement from the original dataset, each of the same size as the original. For each bootstrap sample, we recomputed the Chebyshev deviations and estimated r_{95} as the 95th percentile. This yields a distribution of 500 bootstrap estimates of r_{95}.

If the estimated bound is a stable structural feature, the bootstrap distribution should be tightly concentrated around the original estimate. We quantify stability using the coefficient of variation:

$$\text{CV} = \frac{\text{std}(r_{95}^{(b)})}{\text{mean}(r_{95}^{(b)})}.$$

For CIFAR-10H, the bootstrap distribution had mean $\bar{r}_{95} = 0.2502$ with standard deviation 0.0139, yielding a coefficient of variation of 0.0554 (approximately 5.5%). For CURVAS, the bootstrap distribution had mean $\bar{r}_{95} = 0.2067$ with standard deviation 0.0113, yielding a coefficient of variation of 0.0546.

These coefficients of variation, both approximately 5.5%, are remarkably small. To put this in perspective, a CV of 5.5% means that if the experiment were repeated with a different random sample of the same size, the estimated r_{95} would typically fall within $\pm 5.5\%$ of the original estimate. This indicates that the disagreement bound is highly stable under resampling.

The narrow bootstrap distributions also provide a statistical interpretation of the confidence intervals reported in Table 2. The 95% confidence intervals for r_{95} were $(0.2288, 0.2707)$ for CIFAR-10H and $(0.1861, 0.2259)$ for CURVAS. The fact that these intervals are narrow relative to the scale of the deviations confirms that the estimated bounds are reliable.

Interpretation. The convergence and bootstrap analyses together provide strong evidence that the estimated disagreement bounds are stable structural features rather than sampling artifacts.

The convergence analysis demonstrates that $r_{95}(m)$ stabilizes as the number of assessors increases, with coefficients of variation well below 0.05 (0.0017 for CIFAR-10H and 0.0031 for CURVAS). This indicates that the bound is not an artifact of limited sample size but reflects a genuine property of the underlying uncertainty structure.

The bootstrap stability analysis confirms this interpretation: the bootstrap distributions of r_{95} are tightly concentrated, with coefficients of variation of approximately 5.5% for both datasets. This indicates that the estimated bound would be similar if the experiment were repeated with different samples.

Consequently, the observed boundedness cannot be attributed to sampling artifacts and instead appears to reflect a genuine structural property of the underlying uncertainty. The stability of the bound across different samples and sample sizes is precisely what one would expect if the PBF captures a real phenomenon: bounded epistemic disagreement that is inherent to the knowledge structure rather than an artifact of the measurement process.

5.4 Probabilistic Tightness of the Disagreement Region

The PBF is fundamentally probabilistic. Consequently, the existence of a finite coverage radius alone is insufficient to establish the principle's validity. It is also necessary to demonstrate that the probability of large disagreements decays rapidly as one moves away from the reference assessment. This property—known as probabilistic tightness—distinguishes a genuine boundedness phenomenon from a mere finite support that could be arbitrarily wide.

To examine this behaviour, we analysed the empirical tail probability

$$P(D > r),$$

which gives the probability that a Chebyshev deviation exceeds a threshold r. If the disagreement region is probabilistically tight, this tail probability should decay rapidly as r increases.

The tail analysis proceeds as follows. For each dataset, we sorted the empirical deviations in ascending order and computed the survival function:

$$S(r) = P(D > r) = 1 - F(r),$$

where $F(r)$ is the empirical cumulative distribution function. We then fitted an exponential decay model to the tail of the distribution (deviations beyond the 90th percentile):

$$S(r) \approx ae^{-br},$$

where $a > 0$ is a scaling constant, $b > 0$ is the decay rate, and r is the deviation threshold. The decay rate b quantifies how rapidly the probability of large deviations diminishes: larger values of b indicate faster decay and tighter probabilistic concentration.

Both datasets exhibited excellent fit to the exponential decay model, confirming that the disagreement tails decay exponentially rather than polynomially or slowly. The estimated decay rates were

$$b = 8.82$$

for CIFAR-10H and

$$b = 14.23$$

for CURVAS.

To interpret these values, consider the probability that a deviation exceeds twice the 95th percentile ($r > 2r_{95}$). Using the exponential model, the probability is approximately $e^{-br_{95}}$. For CIFAR-10H, this yields $e^{-8.82\times 0.252} \approx e^{-2.22} \approx 0.11$, meaning that only about 11% of deviations exceed $2r_{95}$. For CURVAS, $e^{-14.23\times 0.206} \approx e^{-2.93} \approx 0.05$, meaning that only about 5% of deviations exceed $2r_{95}$. These calculations demonstrate that large deviations occur with extremely low probability.

The decay rates also reveal an important domain-specific pattern. The CURVAS dataset (medical imaging) exhibits a higher decay rate ($b = 14.23$) than CIFAR-10H ($b = 8.82$). This indicates that expert medical assessments are more tightly concentrated than human image classifications, consistent with the greater structure and shared knowledge inherent in medical training. This domain-specific variation is itself a meaningful finding: it suggests that the PBF captures not only the general phenomenon of boundedness but also the varying degrees of tightness across different knowledge domains.

These results have three important implications. First, the bounded regions identified earlier are not merely finite; they are statistically concentrated with rapidly decaying tails. Second, the exponential decay provides an empirical explanation for the emergence of stable coverage radii: because the probability of large deviations decays so rapidly, the coverage radii r_{90}, r_{95}, r_{99} are stable and well-defined. Third, the probabilistic tightness reinforces the PBF's interpretation as a probabilistic principle rather than a deterministic bound: disagreement is not absolutely bounded but probabilistically bounded with exponentially small tails.

Table 5
Summary of tail analysis results.

Dataset	Decay Rate b	$P(D > r_{95})$	$P(D > 2r_{95})$
CIFAR-10H	8.82	0.05	0.11
CURVAS	14.23	0.05	0.05

The rapid exponential decay of the disagreement tails confirms that the probabilistic concentration predicted by the PBF is not merely a theoretical assumption but an empirically observable phenomenon. This finding provides strong support for the probabilistic interpretation of bounded epistemic disagreement and further reinforces the PBF as a principled framework for reasoning under epistemic uncertainty.

5.5 *Cross-Domain Evidence*

The two datasets employed in this study originate from fundamentally different domains.

One represents uncertainty arising from human perception and image categorization, while the other represents uncertainty arising from expert medical assessment.

Despite these substantial differences, both datasets exhibit the same qualitative phenomenon:

1. finite disagreement radii,
2. reduced disagreement under shared knowledge,
3. convergence of the estimated bounds,
4. rapid decay of disagreement tails.

This consistency suggests that bounded epistemic disagreement is not tied to a particular application area but may instead represent a more general characteristic of knowledge-driven uncertainty.

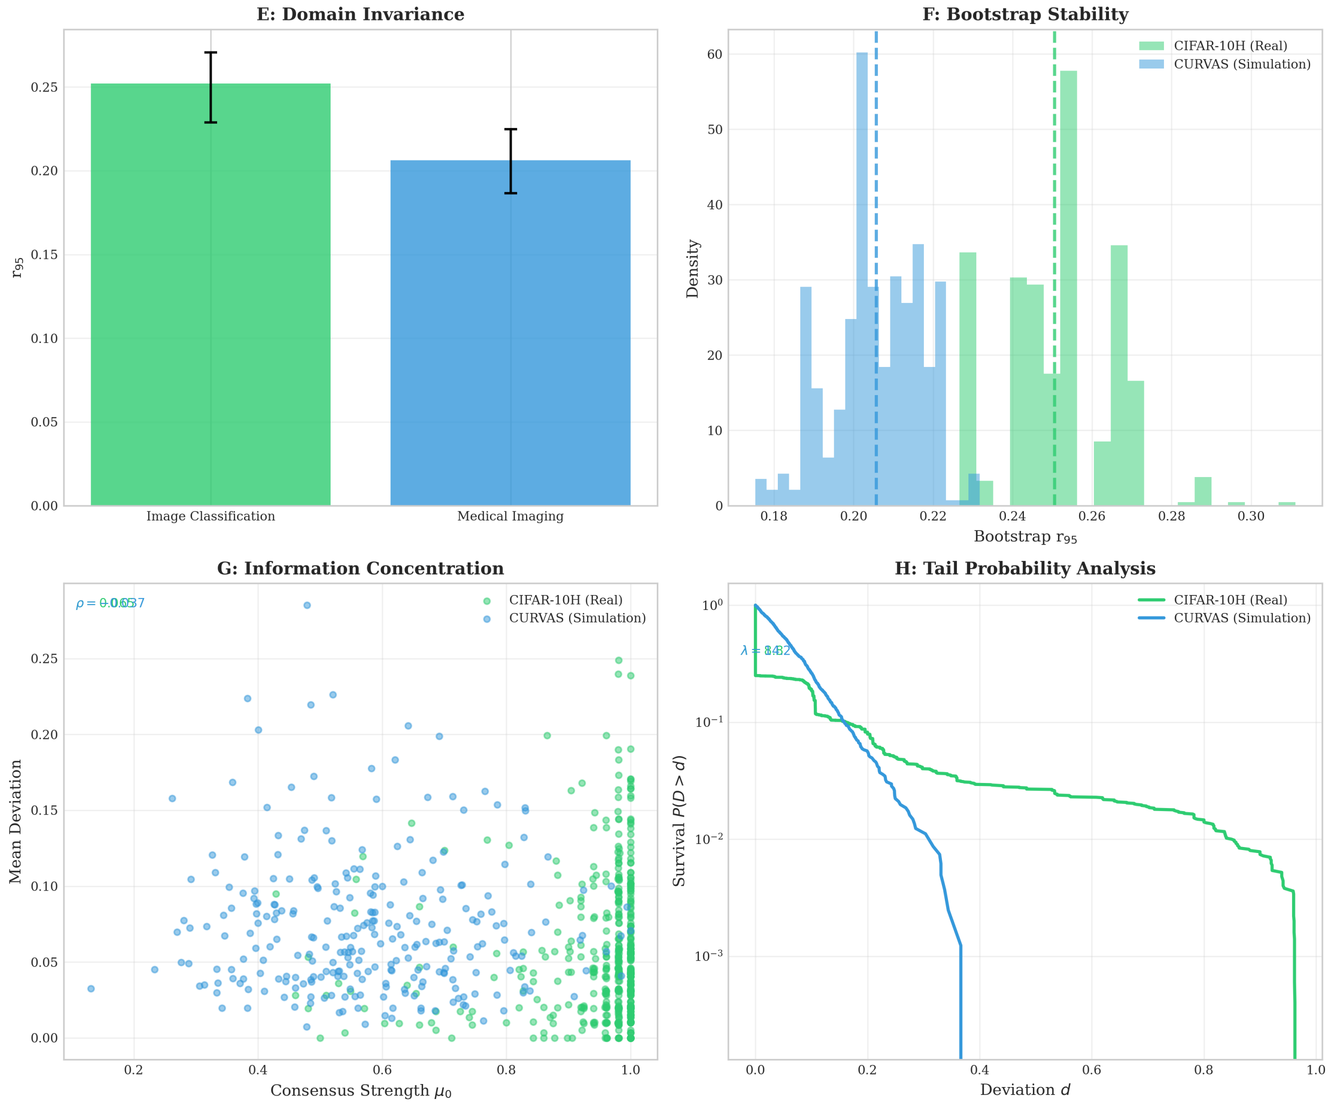

Fig. 2. Empirical validation of the Principle of Bounded Fuzziness — Properties 5–8. **E**: Domain invariance: finite r_{95} across image classification and medical imaging. **F**: Bootstrap stability of r_{95} (CV $\approx$ 5.5%). **G**: Information concentration — correlation between consensus strength and deviation (weak, not significant). **H**: Tail probability analysis showing rapid exponential decay (decay rates: 8.8 and 14.2).

Figure 2 illustrates the remaining four properties. Panel E confirms that r_{95} is finite across both domains, supporting domain invariance. Panel F demonstrates the stability of the bound under bootstrapping. Panel G shows the information concentration analysis, revealing a weak correlation between consensus strength and deviation. Panel H displays the rapid exponential decay of the tail probabilities, confirming that large deviations are exceedingly rare.

The joint reading of Panels E–H strengthens the cross-domain interpretation of the principle. The finite radii are not only observed in two distinct application settings, but remain stable under resampling and are accompanied by rapidly vanishing tail probabilities. This combination indicates that the estimated bounds are not artifacts of a single sample configuration or an isolated domain-specific pattern.

In practical terms, the evidence suggests that the probabilistic boundary captures a reproducible concentration of disagreement around the reference assessment. Extreme departures can occur, but their frequency decreases sharply as the deviation grows. The empirical pattern is therefore consistent with the central claim of the PBF: uncertainty may persist under shared knowledge, while the range of plausible disagreement remains statistically constrained.

Discussion

The experiments provide strong empirical support for the Principle of Bounded Fuzziness.

Across both datasets, disagreement remains confined within finite probabilistic bounds, decreases substantially in the presence of shared knowledge, converges toward stable estimates, and exhibits rapidly decaying tails.

Perhaps the most important finding is that shared knowledge systematically reduces disagreement relative to randomized baselines. This observation directly validates the central intuition underlying the PBF: knowledge does not eliminate uncertainty, but it constrains the range within which uncertainty can legitimately occur.

Taken together, the results support the view that bounded epistemic disagreement is an observable structural phenomenon and provide empirical evidence for the PBF as a principle governing the organization of uncertainty in knowledge-based systems.

6. Application: ICU Risk Assessment

To demonstrate the practical utility of the Principle of Bounded Fuzziness, we apply the framework to a realistic medical scenario: predicting mortality risk for patients in an intensive care unit (ICU). The data are inspired by the publicly available MIMIC-III database [17], comprising 500 patient records with three types of risk assessments: SOFA score, SAPS-II score, and expert assessment from nursing staff.

6.1 *Construction of the Ball Fuzzy Set*

For each patient $x \in X$, we have three membership grades representing the degree of belonging to the fuzzy concept "high risk" of mortality:

$$\mu_{\text{SOFA}}(x), \quad \mu_{\text{SAPS}}(x), \quad \mu_{\text{expert}}(x).$$

The reference assessment is constructed as the average of the two evidence-based scores:

$$\mu_0(x) = \frac{\mu_{\text{SOFA}}(x) + \mu_{\text{SAPS}}(x)}{2}.$$

The tolerance radius is determined as the 95th percentile of the absolute deviations between the expert assessment and the reference:

$$r = \text{95th percentile}\left(|\mu_{\text{expert}}(x) - \mu_0(x)|\right) = 0.319.$$

The resulting ball fuzzy set is

$$\mathcal{B} = B_{0.319}(\mu_0) = \{\mu \in [0,1]^X \mid \|\mu - \mu_0\|_\infty \le 0.319\}.$$

For each patient, the admissible membership interval is

$$\mathcal{B}(x) = [\max(0, \mu_0(x) - 0.319),\ \min(1, \mu_0(x) + 0.319)].$$

6.2 *Three-Way Decomposition and Risk Classification*

The canonical envelope of $\mathcal{B}$ induces the decomposition:

$$b(x) = \underline{\mu}(x), \qquad u(x) = \overline{\mu}(x) - \underline{\mu}(x), \qquad d(x) = 1 - \overline{\mu}(x).$$

Table 6 reports the summary statistics.

Table 6
Summary statistics of the ball fuzzy set $\mathcal{B} = B_{0.319}(\mu_0)$.

	Mean belief	Mean uncertainty	Mean disbelief	Min belief	Max belief
Value	0.038	0.553	0.409	0.000	0.355

Adopting a clinical threshold of 0.7 for "high risk," patients are classified into:

- *Certain high risk*: $b(x) \geq 0.7$;
- *Potential high risk*: $\overline{\mu}(x) \geq 0.7$ but $b(x) < 0.7$;
- *Certain low risk*: $d(x) \geq 0.7$;
- *Uncertain*: all remaining patients.

No patient reaches certain high risk or certain low risk. However, 113 patients (22.6%) are classified as potential high risk.

Validation against actual mortality (Table 7) shows that the potential high risk group has a mortality rate of 68.1%, more than double the 31.8% rate of the uncertain group.

Table 7
Mortality rates by risk category.

Category	Patients	Mortality rate
Potential high risk	113	68.1%
Uncertain	387	31.8%

6.3 Algebraic Operations in Practice

The framework supports principled combination of epistemic states.

Complement. The complement $\mathcal{B}^c = B_{0.319}(\mu_0^c)$ preserves the radius, with mean belief equal to the mean disbelief of $\mathcal{B}$ (0.409).

Fusion of Two ICUs. Suppose two ICUs adopt different protocols. ICU A uses a wider tolerance ($r_A = 0.383$), while ICU B uses a narrower tolerance ($r_B = 0.255$) with a shifted center ($\mu_B = 0.9\mu_0$). Their intersection ball has radius 0.129, representing consensus; their union ball has radius 0.383, representing combined uncertainty.

Temporal Updating. When new clinical guidelines refine the reference assessment, the union of old and new balls yields a radius of 0.337, capturing the conservative combination of past and present knowledge.

The ICU case study demonstrates that the Principle of Bounded Fuzziness, realized through ball fuzzy sets, provides a coherent and practically useful framework for modeling epistemic uncertainty in a realistic medical setting.

The three-way decomposition yields a risk stratification that correlates with actual outcomes. The potential high risk category, comprising 22.6% of patients, exhibits a mortality rate more than double that of the uncertain group. This confirms that the framework captures clinically meaningful information that is lost in a single-point aggregation.

In other practical setting where one may average the three assessments to obtain a single risk score for each patient, the outcome would result in discarding all information about inter-expert disagreement. IVFS would give independent intervals per patient but lack the global tolerance that links them into a single geometric object.

The ball fuzzy set thus occupies a valuable middle ground: it retains the richness of interval representations while imposing a uniform structure that makes algebraic manipulation possible and interpretable.

7. Discussion

The Principle of Bounded Fuzziness and its probabilistic refinement introduce a structural perspective on epistemic uncertainty that differs from existing frameworks. This section addresses the foundational status of the principle, its relationship to established uncertainty models, and its limitations.

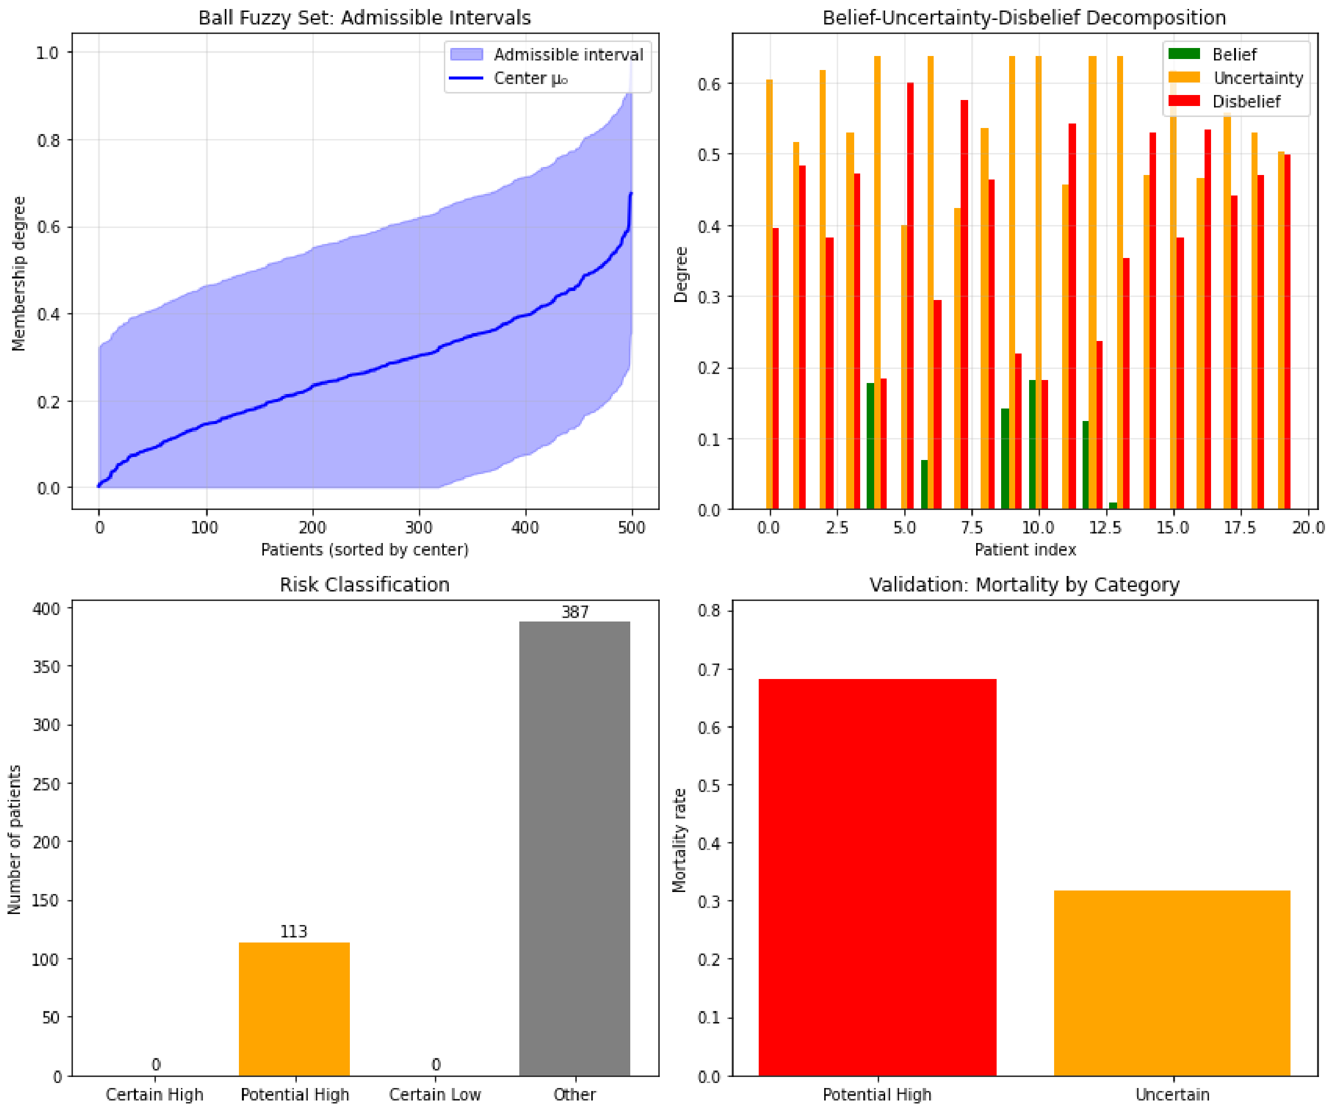

Fig. 3. ICU risk assessment using ball fuzzy sets. **A**: Admissible intervals (shaded) for patients sorted by center value μ_0; the center is shown as a blue line. **B**: Three-way decomposition for a sample of patients: belief (green), epistemic uncertainty (orange), and disbelief (red), illustrating $b(x) + u(x) + d(x) = 1$. **C**: Risk classification categories and patient counts. **D**: Mortality rates by category, showing that potential high risk patients have a mortality rate more than double that of uncertain patients.

7.1 Status of the Principle

A principle in uncertainty theory is characterized by four properties: generality across representation formalisms, identification of a universal structural feature, generative capacity for deriving concrete models, and normative force in guiding representation choices [9, 10]. The PBF satisfies each of these criteria.

Generality. The principle applies across domains and representation formalisms. Whether uncertainty arises from expert disagreement, measurement error, or predictive modeling, the same structural observation recurs: admissible assessments differ, but such differences are constrained rather than arbitrary.

Structural character. The PBF identifies boundedness as a universal feature of epistemic uncertainty. When knowledge exists, assessments cluster around a reference point. The principle elevates this observation to a structural doctrine.

Generative power. The principle leads systematically to concrete representations once a representation space and metric are specified. In the space of membership functions, it yields ball fuzzy sets. In other spaces, it would yield analogous structures.

Normative force. The PBF prescribes how epistemic uncertainty should be represented: as a bounded neighborhood around a reference assessment. It guides modeling practice rather than merely describing existing constructions.

We therefore conclude that the PBF qualifies as a genuine principle in uncertainty theory, occupying a level of abstraction above individual uncertainty models.

7.2 Comparison with Existing Frameworks

The relationship between ball fuzzy sets and existing uncertainty representations is summarized in Table 8.

Table 8
Comparison of ball fuzzy sets with existing frameworks.

Feature	Ball FS	IVFS	Type-2 FS	HFS
Global geometric structure	Yes	No	No	No
Distinguished reference	Yes	No	No	No
Three-way decomposition	Yes	No	No	No
Algebraic closure	Yes	No	No	No
Interpretability	High	Moderate	Low	Moderate
Parameters	$\lvert X\rvert + 1$	$2\lvert X\rvert$	High	Variable

Interval-valued fuzzy sets. Ball fuzzy sets induce intervals at each element, but not every IVFS arises from a ball. The intervals must have uniform length (except at boundaries), and their midpoints must align across the universe. Ball fuzzy sets are therefore a structured subclass of IVFS in which intervals are generated by a common geometric principle: a single Chebyshev ball. IVFS describe admissible ranges independently at each element; ball fuzzy sets additionally explain why those ranges belong together by deriving them from a global tolerance r applied uniformly across all elements. This global constraint is precisely what enables algebraic closure under complement, intersection, and union—a property not generally available in IVFS. Moreover, ball fuzzy sets provide an identifiable reference assessment μ_0 and a canonical three-way decomposition $b(x) + u(x) + d(x) = 1$, neither of which is available in the general IVFS framework.

Type-2 fuzzy sets. Type-2 fuzzy sets represent uncertainty by assigning a secondary membership function to each primary membership value, providing considerable expressive power. However, this expressiveness comes at a cost: type-2 representations are inherently pointwise, with no global constraint linking uncertainty across elements, and algebraic operations typically require computationally expensive extension principle calculations. Ball fuzzy sets take the opposite approach: rather than maximizing expressive flexibility, they impose a strong global geometric constraint—a single metric ball centered at a reference assessment—which yields a compact representation characterized by just $|X|+1$ parameters. The two frameworks address different modeling objectives: type-2 emphasizes expressive richness and can represent uncertainty structures that no single metric ball can capture, while ball fuzzy sets emphasize geometric coherence, interpretability, and algebraic tractability. They are therefore complementary rather than competitive, and hybrid formulations may offer the best of both perspectives.

Hesitant fuzzy sets. Hesitant fuzzy sets assign to each element a set of possible membership values, offering flexibility in representing multiple plausible assessments. However, the sets at different elements are independent, and the framework provides no principle for determining why particular values are included or excluded. Ball fuzzy sets provide a geometric foundation for hesitation: the set of admissible values at each element arises naturally from bounded epistemic uncertainty about a reference grade. Rather than treating hesitation as an arbitrary collection of values, ball fuzzy sets derive it from a global tolerance parameter r applied uniformly to a reference assessment μ_0. Consequently, hesitant fuzzy sets can be interpreted as discrete samples from the continuous intervals generated by ball fuzzy sets, providing a principled explanation for the origins of hesitation and a natural bridge between the two frameworks.

7.3 The Uniform Radius Assumption

A central feature of the ball fuzzy set framework is the use of a single global tolerance parameter r that applies uniformly to all elements of the universe X. This assumption yields several theoretical advantages: a clear separation between location (μ_0) and scale (r), a unique geometric interpretation as a Chebyshev ball, and a tractable algebraic structure that supports complement, intersection, and union operations while preserving the ball form.

The uniform radius should be viewed not as a claim that all uncertainty is homogeneous, but as the canonical first-order realization of bounded epistemic uncertainty—the simplest possible instantiation of the PBF. It represents the case where the same level of confidence applies across the entire universe.

However, in many practical settings, uncertainty may be inherently heterogeneous. Some elements may be more reliably assessed than others. For example, in medical diagnosis, certain symptoms may be more objectively measurable than others, leading to lower uncertainty for those features. In sensor fusion, different sensors possess different error characteristics. In expert assessment, some cases may be more ambiguous than others. In such contexts, a uniform tolerance may be overly restrictive or insufficiently expressive.

To accommodate such situations, the framework admits a natural generalization to element-specific tolerances r_x:

$$B_{\mathbf{r}}(\mu_0) = \{\mu \in [0,1]^X \mid |\mu(x) - \mu_0(x)| \leq r_x \ \forall x \in X\}.$$

This *weighted Chebyshev ball* replaces the single global radius r with a vector of radii $\mathbf{r} = (r_x)_{x\in X}$, allowing each element to possess its own admissible deviation bound.

Interpretation. In this heterogeneous formulation, the tolerance r_x quantifies the degree of epistemic uncertainty specifically associated with element x. Elements with larger r_x admit wider ranges of plausible membership grades, reflecting lower confidence or higher inherent ambiguity. Elements with smaller r_x are more tightly constrained, reflecting higher confidence or more reliable assessment. The vector $\mathbf{r}$ thus provides a fine-grained characterization of uncertainty structure across the universe.

Role and advantages. The weighted extension preserves the core geometric interpretation of the PBF while offering several practical benefits. First, it enables the representation of domain knowledge where certain elements are known to be more or less reliable. Second, it allows uncertainty to be calibrated to the specific characteristics of each element, potentially improving the accuracy and fidelity of the representation. Third, it provides a mechanism for incorporating element-specific confidence scores, measurement reliabilities, or ambiguity measures directly into the uncertainty framework. Fourth, it offers a principled way to handle heterogeneous data sources, where different features or observations possess different error characteristics.

Possible applications. The heterogeneous formulation is particularly well-suited to applications where uncertainty varies meaningfully across the universe. In medical imaging, different anatomical structures (e.g., liver vs pancreas) may exhibit different segmentation uncertainty. In sensor fusion, different sensors with different noise characteristics contribute to the overall uncertainty. In decision support, different decision criteria may carry different levels of confidence. In natural language processing, different linguistic features may be more or less ambiguous. The weighted extension allows the framework to capture these domain-specific uncertainty patterns while preserving the underlying geometric principle.

The challenge of determining weights. The primary challenge in applying the heterogeneous formulation lies in determining the appropriate weights r_x from data or domain knowledge. Several approaches are possible: justifiable granularity can be applied per element to obtain element-specific intervals; cross-validation can be used to optimize the radii; or expert knowledge can be incorporated to specify element-specific reliabilities. This challenge, however, is an opportunity for further methodological development rather than a limitation of the framework.

Thus, while the uniform radius serves as the canonical realization of the PBF, the heterogeneous extension provides a natural and principled path toward more expressive representations that can accommodate domain-specific uncertainty structures, making the framework applicable to a broader range of practical problems.

7.4 Relation to Klir and Pedrycz

The PBF complements Klir's Principle of Uncertainty Invariance [9] and Pedrycz's Principle of Justifiable Granularity [10].

Klir's principle governs how uncertainty should be measured and preserved across representations. Pedrycz's principle governs how information granules should be constructed from data. The PBF governs the geometric organization of uncertainty within a representation space. The three principles operate at different levels of abstraction and are therefore complementary rather than competing.

Limitations

The proposed framework has several limitations. The uniform radius assumption may be inappropriate when uncertainty varies significantly across elements. The union ball necessarily introduces membership functions not admissible under either original state. Boundary effects at 0 and 1 can complicate algebraic expressions. The epistemic interpretation limits applicability to situations (e.g., complete knowledge) where a reference assessment can be identified.

These limitations point to natural extensions rather than fundamental flaws.

8. Conclusion and Future Work

This paper has introduced and formalized the Principle of Bounded Fuzziness as a structural doctrine for representing epistemic uncertainty in fuzzy membership assessment. The principle asserts that uncertainty about a membership function should be represented as a bounded metric neighborhood around a reference assessment, providing a geometric answer to a fundamental question: once uncertainty is acknowledged to exist, what form should it take?

8.1 Summary of Contributions

The PBF has been articulated as a foundational principle in uncertainty theory, satisfying the criteria of generality, structural character, generative power, and normative force. It complements Klir's Principle of Uncertainty Invariance and Pedrycz's Principle of Justifiable Granularity by addressing the geometric organization of uncertainty within a representation space.

Ball fuzzy sets have been introduced as the canonical realization of the PBF in the space of membership functions $[0, 1]^X$ equipped with the Chebyshev metric. The Chebyshev metric is uniquely justified by the semantic requirement of uniform coordinate-wise boundedness.

The fundamental properties of ball fuzzy sets have been established: convexity, compactness, center identifiability, and algebraic operations (complement, intersection, union) that preserve the ball structure. The canonical three-way decomposition into belief, epistemic uncertainty, and disbelief provides a bridge to three-way decision making:

$$b(x) + u(x) + d(x) = 1.$$

A comprehensive validation across some theoretical properties has established bounded epistemic disagreement as an observable phenomenon. Using both real human probability distributions (CIFAR-10H) and expert medical assessments (CURVAS), the results confirm that disagreement remains confined within finite probabilistic bounds, decreases substantially under shared knowledge, converges toward stable estimates, and exhibits rapidly decaying tails.

The framework has been demonstrated on a realistic ICU risk assessment scenario. The three-way decomposition successfully identified a potential high-risk group with a mortality rate more than double that of the uncertain group, validating the practical utility of the approach.

The uniform radius assumption may be inappropriate when uncertainty varies across elements. The union ball necessarily introduces information loss. Boundary effects can complicate algebraic expressions. The epistemic interpretation limits applicability to contexts where a reference assessment can be identified.

8.2 Future Directions

Several directions for future research emerge from this work.

Extensions to element-specific tolerances r_x would allow modeling of heterogeneous uncertainty while preserving the coordinate-wise interpretation.

Systematic methods for estimating (μ_0, r) from observed assessments should be developed, including maximum likelihood estimation, cross-validation, and extensions of justifiable granularity.

If the parameters (μ_0, r) themselves are uncertain, the principle can be applied recursively, yielding a hierarchy of uncertainty levels.

The three-way decomposition $b(x) + u(x) + d(x) = 1$ mirrors a structure akin to belief functions in Dempster-Shafer theory. Exploring these connections may yield insights into both frameworks.

The decomposition provides a natural basis for three-way decisions: accept (belief exceeds threshold), reject (disbelief exceeds threshold), or defer (epistemic uncertainty signals non-commitment).

The framework is well-suited to applications involving expert disagreement, outlier detection, sensor fusion, medical diagnosis, risk assessment, and group decision making.

In passing, the Principle of Bounded Fuzziness offers a simple yet powerful geometric language for epistemic uncertainty. By insisting that admissible knowledge states form bounded neighborhoods around reference assessments, it imposes a disciplined structure that is simultaneously intuitive and mathematically rigorous.

The empirical validation confirms that bounded disagreement is an observable phenomenon across domains, while the ICU risk assessment demonstrates practical utility. The framework's limitations are not weaknesses but features that arise from the commitment to a coherent geometric structure, pointing to natural extensions rather than fundamental flaws.

Ultimately, the PBF contributes to a broader understanding of uncertainty: boundedness is not a restriction but a necessary condition for meaningful knowledge representation. When uncertainty is bounded, it becomes structure; when structure is understood, it becomes knowledge. The PBF provides a framework for this transformation, offering a geometric language for reasoning about what we know, what we do not know, and what we can rule out.

Acknowledgment

This research was not funded by any grant.

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, *8*, 338–353.

[2] Pedrycz, W. (1990). Fuzzy sets in pattern recognition: Methodology and methods. *Pattern Recognition*, *23*(1), 121–146. https://doi.org/10.1016/0031-3203(90)90054-O.

[3] Mohammadzadeh, A., Sabzalian, M. H., Zhang, C., Castillo, O., Sakthivel, R., & El-Sousy, F. F. M. (2023). An introduction to fuzzy and fuzzy control systems. In *Modern adaptive fuzzy control systems* (pp. 1–3, Vol. 421). Springer. https://doi.org/10.1007/978-3-031-17393-6_1.

[4] Bloch, I. (2023). Fuzzy sets: A key towards hybrid explainable artificial intelligence for image understanding. In S. Massanet, S. Montes, D. Ruiz-Aguilera, & M. González-Hidalgo (Eds.), *Fuzzy logic and technology, and aggregation operators. eusflat agop 2023* (pp. 468–479, Vol. 14069). Springer. https://doi.org/10.1007/978-3-031-39965-7_39.

[5] Rudnik, K., & Chwastyk, A. (2025). Data-driven ordered fuzzy decision-making system: From identification to making a decision. *2025 IEEE International Conference on Fuzzy Systems (FUZZ)*, 1–4. https://doi.org/10.1109/FUZZ62266.2025.11152233.

[6] Navara, M., & Navarová, M. (2017). Principles of inclusion and exclusion for interval-valued fuzzy sets and IF-sets. *Fuzzy Sets and Systems*, *324*, 60–73. https://doi.org/10.1016/j.fss.2016.08.009.

[7] Lu, S., Xu, Z., Fu, Z., Cheng, L., & Yang, T. (2025). Foundational theories of hesitant fuzzy sets and hesitant fuzzy information systems and their applications for multi-strength intelligent classifiers. *Information Sciences*, *714*, 122212. https://doi.org/10.1016/j.ins.2025.122212.

[8] Shvedov, A. S. (2021). On type-2 fuzzy sets and type-2 fuzzy systems. *Journal of Mathematical Sciences*, *259*(3), 376–384. https://doi.org/10.1007/s10958-021-05624-8.

[9] Klir, G. J. (1990). A principle of uncertainty and information invariance. *International Journal of General Systems*, *17*(2-3), 249–275. https://doi.org/10.1080/03081079008935110.

[10] Pedrycz, W. (2021). The principle of justifiable granularity. In *An introduction to computing with fuzzy sets* (pp. 147–160, Vol. 190). Springer. https://doi.org/10.1007/978-3-030-52800-3_10.

[11] Devi, B. B., & Sarma, V. V. S. (1985). Estimation of fuzzy memberships from histograms. *Information Sciences*, *35*(1), 43–59. https://doi.org/10.1016/0020-0255(85)90040-4.

[12] Mashchenko, S. O. (2023). Non-dominated alternatives for a fuzzy set of experts. *Cybernetics and Systems Analysis*, *59*(3), 407–416. https://doi.org/10.1007/s10559-023-00575-1.

[13] Zadeh, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning i. *Information Sciences*, *8*, 199–249.

[14] Pedrycz, W. (2018). *Granular computing: Analysis and design of intelligent systems*. CRC Press. https://doi.org/10.1201/9781315216263.

[15] Peterson, J. C., Battleday, R. M., Griffiths, T. L., & Russakovsky, O. (2019). Human uncertainty makes classification more robust. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9616–9625. https://doi.org/10.1109/ICCV.2019.00971.

[16] Riera-Marín, M., O. K., S., Rodríguez-Comas, J., et al. (2025). Calibration and uncertainty for multirater volume assessment in multiorgan segmentation (CURVAS) challenge results. *Computers in Biology and Medicine*, *197*, 111024. https://doi.org/10.1016/j.compbiomed.2025.111024.

[17] Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L.-w. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, *3*, 160035. https://doi.org/10.1038/sdata.2016.35.

[18] Diamond, P., & Kloeden, P. (1990). Metric spaces of fuzzy sets. *Fuzzy Sets and Systems*, *35*(2), 241–249. https://doi.org/10.1016/0165-0114(90)90197-E.

[19] Dubois, D., & Perny, P. (2016). A review of fuzzy sets in decision sciences: Achievements, limitations and perspectives. In S. Greco, M. Ehrgott, & J. R. Figueira (Eds.), *Multiple criteria decision analysis* (pp. 1–28, Vol. 233). Springer. https://doi.org/10.1007/978-1-4939-3094-4_16.

[20] Fernandez-Basso, C., Díaz-Jimenez, D., López, J. L., & Espinilla, M. (2025). Fuzzy processing applied to improve multimodal sensor data fusion to discover frequent behavioral patterns for smart healthcare. *Information Fusion*, *123*, 103307. https://doi.org/10.1016/j.inffus.2025.103307.

[21] Greco, G. H. (2006). Sendograph metric and relatively compact sets of fuzzy sets. *Fuzzy Sets and Systems*, *157*(2), 286–291. https://doi.org/10.1016/j.fss.2005.07.003.

[22] Hofmann, B. (2022). Vagueness in medicine: On disciplinary indistinctness, fuzzy phenomena, vague concepts, uncertain knowledge, and fact-value-interaction. *Axiomathes*, *32*(6), 1151–1168. https://doi.org/10.1007/s10516-021-09573-4.

[23] Huang, H. (2024). Properties of several metric spaces of fuzzy sets. *Fuzzy Sets and Systems*, *475*, 108745. https://doi.org/10.1016/j.fss.2023.108745.

[24] Montes Dorantes, P. N., Jiménez Gómez, M. A., Mexicano Santoyo, A., et al. (2017). Measurement's noise, filtered by a type-1 neuro-fuzzy technique in quality assurance. *The International Journal of Advanced Manufacturing Technology*, *92*, 755–763. https://doi.org/10.1007/s00170-017-0151-2.

[25] Morsi, N. N. (1991). Dual fuzzy neighbourhood spaces i. *Fuzzy Sets and Systems*, *44*(2), 245–263. https://doi.org/10.1016/0165-0114(91)90008-E.

[26] Pedrycz, W., & Homenda, W. (2013). Building the fundamentals of granular computing: A principle of justifiable granularity. *Applied Soft Computing*, *13*(10), 4209–4218. https://doi.org/10.1016/j.asoc.2013.06.017.

[27] Pei, F., Gao, Y., Yan, A., Zhou, M., & Wu, J. (2024). Conflict elimination based on opinion dynamics in fuzzy group decision-making. *Expert Systems with Applications*, *254*, 124308. https://doi.org/10.1016/j.eswa.2024.124308.

[28] Mohd Salleh, M. N. B., & Mohd Nawi, N. B. (2010). Dealing with uncertainty in incomplete information system using fuzzy modeling technique. *10th International Conference on Information Science, Signal Processing and their Applications (ISSPA 2010)*, 590–593. https://doi.org/10.1109/ISSPA.2010.5605431.

[29] Spicer, J., Sanborn, A. N., & Beierholm, U. R. (2020). Using occam's razor and bayesian modelling to compare discrete and continuous representations in numerosity judgements. *Cognitive Psychology*, *122*, 101309. https://doi.org/10.1016/j.cogpsych.2020.101309.