A Category of Blur and the Grand Lemma

Aleksandar Perišić

August 2025

Abstract

We axiomatize *blur*—a tunable neighbourhood/relaxation—as a categorical construction. A blurred object is a diagram $\mathcal{B}_X : I \to \mathcal{C}$ indexed by blur scales with a *reading* map $\rho_X : \lim_I \mathcal{B}_X \to X$ (sharp limit) and exhaustion $\mathrm{colim}_I \mathcal{B}_X \simeq 1$. An uncertainty window is a lax retract pair $j_X^\varepsilon : X \to B_\varepsilon X$, $r_X^\varepsilon : B_\varepsilon X \to X$. We prove a *property-transport* principle: for properties stable under retracts and filtered limits and monotone for $\precsim$, $\mathsf{P}(X)$ holds iff $\mathsf{P}_\varepsilon(B_\varepsilon X)$ holds eventually. Blur-morphisms are natural transformations; the *Grand Lemma (Blur–transport lemma)* shows sending sharp equals sending blurred then reading, yielding a simple string-diagram calculus and a graded (co)monad structure for scale composition. Examples include analytic blur (Markov/convolution semigroups) and logical blur (quantale-enriched nuclei), linking the construction to domain-theoretic approximation and probabilistic powerdomains.

We fix a base category $\mathcal{C}$ (e.g. **Set**, **Top**, or **Meas**). Intuitively, objects are *proposition spaces* (or answer spaces), and a *blur* is a tunable neighbourhood/relaxation.

Blur index and blurred objects

Let $(I, \preceq, \oplus, 0, \infty)$ be a directed monoidal poset of blur scales (think: $\varepsilon \in (0, \infty]$, with $0 =$ sharp, $\infty =$ "no response"), where $\oplus$ models composing blur (e.g. variance addition), 0 is the neutral element, and ∞ is absorbing ($\varepsilon \oplus \infty = \infty$).

Definition 1 (Blurred object)**.** A *blurred object* is a pair $\mathbf{X} = (X, \mathcal{B}_X)$ where $X \in \mathrm{Ob}(\mathcal{C})$ and $\mathcal{B}_X : I \to \mathcal{C}$ is a functor (a filtered diagram) with:

1. **Reading cone (sharp limit):** a universal cone $\rho_X : \lim_I \mathcal{B}_X \to X$. We say the blur is *faithful* if ρ_X is an isomorphism (i.e. $X \simeq \lim_I \mathcal{B}_X$).
2. **Exhaustion at infinity:** a cocone to a terminal/"no-info" object 1, i.e. $\mathrm{colim}_I \mathcal{B}_X \simeq 1$ ("neighbourhood tends to no response").

Write $B_\varepsilon X$ for $\mathcal{B}_X(\varepsilon)$ and $r_X^\varepsilon : B_\varepsilon X \to X$ for the leg of the limit cone (*reading at scale* ε).

Remark 2 (Relation to domain-theoretic approximation)**.** *If we regard I as a directed set of "resolutions", a faithful blur has the flavour of domain theory: X is recovered as the limit of its approximants $(B_\varepsilon X)_{\varepsilon \in I}$. The exhaustion condition plays the role of a bottom element, and the blur diagram is analogous to a directed system of finite approximations whose limit is the full object. The difference is that here the approximation is explicitly parametrized by a physical/epistemic scale ε.*

Remark 3. *Concrete examples:*

- $\mathcal{C} =$ **Meas**, *$B_\varepsilon X = X$ equipped with a Gaussian jitter of variance ε (convolution on functions, or Giry-type randomization on points);*

- $\mathcal{C} = \mathbf{Top}$, *$B_\varepsilon X$ the same set with a coarser uniformity/entourage (neighbourhood thickening);*
- $\mathcal{C} = \mathbf{Set}$, *$B_\varepsilon X$ the set of "ε-consistent truth tables" for X with a collapse map to X as $\varepsilon \to 0$.*

Weak containment and the uncertainty window

The point of blur is that an object is still *present* while travelling with a neighbourhood, but only in a relaxed sense. We capture this by a *lax retract pair.*

Definition 4 (Uncertainty window / lax retract pair)**.** For each scale $\varepsilon \in I$, we equip X with a pair of natural maps

$$j_X^\varepsilon : X \longrightarrow B_\varepsilon X \qquad \text{and} \qquad r_X^\varepsilon : B_\varepsilon X \longrightarrow X,$$

called the *injection (thickening)* and the *reading (deblurring)*, such that

$$r_X^\varepsilon \circ j_X^\varepsilon \;=\; \mathrm{id}_X \qquad \text{and} \qquad j_X^\varepsilon \circ r_X^\varepsilon \;\precsim\; \mathrm{id}_{B_\varepsilon X}.$$

Here $\precsim$ is a chosen ambient preorder on endomorphisms (e.g. pointwise $\leq$ in **Set** or **Top** when available, a.s. $\leq$ in **Meas**, or the enrichment order if $\mathcal{C}$ is Pos/quantale-enriched). Thus X is *contained while blurred* (exact on the sharp side, relaxed on the blurred side). We also require naturality in X and that $j_X^\varepsilon \to \mathrm{id}_X$ and $r_X^\varepsilon \to \mathrm{id}_X$ as $\varepsilon \to 0$ (in the sense of your reading cone).

Remark 5 (Both–sides neighbourhood and synchronized removal)**.** *The pair $(j_X^\varepsilon, r_X^\varepsilon)$ says: as long as the Proposition and the Answer carry* matching *neighbourhoods, the mechanism can pass them with controlled relaxation, and removing the neighbourhood on one side forces the corresponding removal on the other via naturality of $\bar{f}$ and the graded (co)monad maps $B_{\varepsilon\oplus\delta}X \to B_\varepsilon(B_\delta X)$.*

Remark 6 (Transmission principle (AC-analogy))**.** *Think of ε as an* impedance *knob. Sending X through the network at scale ε uses j_X^ε to load X into $B_\varepsilon X$, transports with $\bar{f}_\varepsilon$, then reads with r_Y^ε. Some invariants (mass, expectations, zero-mode, conserved charges) pass* losslessly *because $\bar{f}_\varepsilon$ and r^ε preserve them; sharp, fragile invariants are replaced by their relaxed versions while the signal is AC-carried (blurred). As $\varepsilon \downarrow 0$ the impedance vanishes and the relaxation collapses back to sharp equality.*

Proposition 7 (Property transport under uncertainty)**.** *Let P be a property of objects that is (i) stable under retracts and filtered limits, and (ii) monotone with respect to the ambient preorder $\precsim$ on endomorphisms. Define its ε-relaxation P_ε by "$B_\varepsilon X$ has P up to $\precsim$". Then for faithful blurs:*

$$\mathsf{P}(X) \quad \Longleftrightarrow \quad \big(\exists\, \varepsilon_0\ \forall\, \varepsilon \preceq \varepsilon_0 :\ \mathsf{P}_\varepsilon(B_\varepsilon X)\big).$$

In words: an object has P iff *it has the relaxed property throughout some uncertainty window, and reading removes the relaxation without loss.*

Sketch. ($\Rightarrow$) If $\mathsf{P}(X)$ holds and P is stable under retracts, then $B_\varepsilon X$ inherits P_ε via j_X^ε and r_X^ε (lax retract). ($\Leftarrow$) If $\mathsf{P}_\varepsilon(B_\varepsilon X)$ holds eventually and $\rho_X : \lim_I B_\varepsilon X \xrightarrow{\cong} X$, stability under filtered limits transfers P to X; monotonicity in $\precsim$ removes the laxity in the limit. $\square$

Remark 8 (Blur as a graded comonad)**.** *The family $(B_\varepsilon)_{\varepsilon \in I}$ comes with canonical comparison maps $B_{\varepsilon\oplus\delta}X \to B_\varepsilon(B_\delta X)$, natural in X, making it a* graded comonad *(or dually, a graded monad) on $\mathcal{C}$. In concrete analytic models these maps are realized by Markov semigroups or convolution semigroups; in logical models they correspond to iterating a nucleus or closure operator. Thus the categorical structure behind blur is compatible with standard constructions in probability (Giry/Markov) and in logic (modal/comonadic viewpoints).*

Morphisms that respect blur

Definition 9 (Blur-morphism)**.** Given blurred objects $\mathbf{X} = (X, \mathcal{B}_X)$ and $\mathbf{Y} = (Y, \mathcal{B}_Y)$, a *blur-morphism* $\bar{f} : \mathbf{X} \to \mathbf{Y}$ is a natural transformation

$$\bar{f} : \mathcal{B}_X \Rightarrow \mathcal{B}_Y, \qquad \text{i.e. } \forall \varepsilon \in I \;\; \bar{f}_\varepsilon : B_\varepsilon X \to B_\varepsilon Y \;\text{ and }\; \bar{f}_{\varepsilon \preceq \varepsilon'} \text{ commute with reindexing.}$$

Its *sharp part* is $f := \rho_Y \circ \lim_I \bar{f} \circ \rho_X^{-1} : X \to Y$ whenever the blurs are faithful.

Composition is pointwise: $(\bar{g} \circ \bar{f})_\varepsilon := \bar{g}_\varepsilon \circ \bar{f}_\varepsilon$; identities are $(\mathrm{id})_\varepsilon = \mathrm{id}$. Thus blurred objects/morphisms form a category $\mathbf{Blur}(\mathcal{C})$.

Remark 10 (Bidirectional observability)**.** *A network is* observable in both directions *if there exists $\bar{f}^\dagger : \mathbf{Y} \to \mathbf{X}$ with*

$$\rho_X \circ \lim_I (\bar{f}^\dagger \circ \bar{f}) \circ \rho_X^{-1} = \mathrm{id}_X, \qquad \rho_Y \circ \lim_I (\bar{f} \circ \bar{f}^\dagger) \circ \rho_Y^{-1} = \mathrm{id}_Y,$$

i.e. $\bar{f}$ is an isomorphism in $\mathbf{Blur}(\mathcal{C})$*. This formalizes that we may switch Proposition/Answer at will.*

Unison removal and scale composition

We assume reindexing is monoidal: for each $\varepsilon, \delta \in I$ there is a canonical comparison $\mu_{\varepsilon,\delta} : B_{\varepsilon \oplus \delta} X \to B_\varepsilon(B_\delta X)$, natural in X, making $(B_\varepsilon)_{\varepsilon \in I}$ a *graded comonad* (or graded monad, depending on the concrete model). Naturality of $\bar{f}$ means *removing neighbourhoods happens in unison*:

$$\bar{f}_\varepsilon \;\text{ commutes with the transition maps } B_{\varepsilon'} X \to B_\varepsilon X \;\text{ and }\; B_{\varepsilon'} Y \to B_\varepsilon Y.$$

Standing convention. We fix natural j_X^ε as in Definition 4, so $X \xrightarrow{j_X^\varepsilon} B_\varepsilon X \xrightarrow{r_X^\varepsilon} X$ is a lax retract pair (object contained with an uncertainty window).

Grand Lemma (sequential deblurring = one-shot deblurring)

Theorem 11 (Grand Lemma (Blur–transport lemma): sending sharp ⇔ sending blurred and reading stepwise)**.** *Let* $\mathbf{X} \xrightarrow{\bar{f}} \mathbf{Y} \xrightarrow{\bar{g}} \mathbf{Z}$ *be blur-morphisms with faithful blurs. Then for every $\varepsilon \in I$,*

$$r_Z^\varepsilon \circ (\bar{g}_\varepsilon \circ \bar{f}_\varepsilon) \;=\; r_Z^\varepsilon \circ \bar{g}_\varepsilon \circ r_Y^\varepsilon \circ \bar{f}_\varepsilon \;=\; (g \circ f) \circ r_X^\varepsilon,$$

and in particular, after taking the limit $\varepsilon \to 0$ (reading),

$$g \circ f \;=\; \rho_Z \circ \lim_I (\bar{g} \circ \bar{f}) \circ \rho_X^{-1} \;=\; \rho_Z \circ \lim_I \bar{g} \circ \rho_Y^{-1} \circ \rho_Y \circ \lim_I \bar{f} \circ \rho_X^{-1}.$$

Thus sending a Proposition through the network is equivalent to sending it together with its neighbourhood and reading in successive steps.

Proof. Pointwise identity $r_Z^\varepsilon \circ \bar{g}_\varepsilon \circ \bar{f}_\varepsilon = (g \circ f) \circ r_X^\varepsilon$ follows from naturality of the limit cones: r_Y^ε and r_Z^ε are the legs of ρ_Y, ρ_Z, and $\bar{f}, \bar{g}$ are natural transformations. Taking limits in I yields the equalities with ρ_X, ρ_Y, ρ_Z by the universal property of limits. □

Reading at infinity (no response)

If $\mathrm{colim}_I \mathcal{B}_X \simeq 1$ and $\mathrm{colim}_I \mathcal{B}_Y \simeq 1$, then any blur-morphism $\bar{f}$ induces the unique arrow $1 \to 1$ at ∞; i.e. as neighbourhood grows without bound both Proposition and Answer collapse to "no response" in unison.

Passing neighbourhoods faithfully

Given any neighbourhood ε on the Proposition, the network passes it to the Answer via $\bar{f}_\varepsilon$; the family $\{\bar{f}_\varepsilon\}_{\varepsilon \in I}$ ensures the Answer retains the *corresponding* neighbourhood. Rescaling of neighbourhoods is handled by reindexing I (e.g. $\varepsilon \mapsto \alpha\varepsilon$) and the graded (co)monad coherence.

Two canonical models

1. **Analytic blur (convolution).** $\mathcal{C} = \mathbf{Meas}$, $B_\varepsilon X$ acts on measurable functions by convolution with a positive kernel K_ε (Gaussian/Poisson/compact mollifier), and on points via randomization with law K_ε. Then $\bar{f}$ is a Markov kernel family commuting with convolution; ρ is the deconvolution limit as $\varepsilon \to 0$.

2. **Logical blur (fuzzy/enriched).** $\mathcal{C}$ enriched over the quantale $([0,1], \leq, \cdot, 1)$: a proposition has a neighbourhood of truth values; B_ε thickens truth via a nucleus (closure operator). Morphisms are $[0,1]$-nonexpansive maps. Limits recover crisp truth; colimits at ∞ give the trivial truth.

Remark 12 (Stochastic universality)**.** *Instead of fixed kernels, let U_ε be random perturbations with $\mathbb{P}(|U_\varepsilon| > \delta) \to 0$ as $\varepsilon \to 0$. Then $B_\varepsilon f(x) := \mathbb{E}[f(x + U_\varepsilon)]$ defines a blurred diagram; all statements above hold verbatim. "Gaussian" is merely a maximally symmetric instance.*

String-diagram view of blur

We now depict blur morphisms as string diagrams. Wires are objects, small boxes are processes, triangles are *reading* (removing the blur).

Symbols

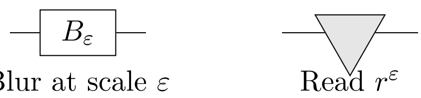

Grand Lemma as a diagram

Sending a proposition X through a network f and then reading is equal to blurring first, sending through the blurred network $\bar{f}$, then reading. This is the commutativity of the diagram:

$$\begin{tikzcd} X \arrow[r, "r^\varepsilon_X"] \arrow[rrr, bend left, "f"] & B_\varepsilon X \arrow[r, "\bar{f}_\varepsilon"] & B_\varepsilon Y \arrow[r, "r^\varepsilon_Y"] & Y \end{tikzcd}$$

This string-diagram encapsulates the *Grand Lemma (Blur–transport lemma)*:

$$r^\varepsilon_Y \circ \bar{f}_\varepsilon \circ r^\varepsilon_X \quad = \quad f \circ r^\varepsilon_X.$$

Bidirectionality

If $\bar{f}$ is an isomorphism in the blur category, the diagram works both ways: we may slide the reading triangle across the network in either direction, expressing the *observability in both roles* of Proposition and Answer.

Conclusion: The Grand Purpose of Blurring

The central role of blur is not merely technical but conceptual.

Grand purpose. Blurring allows an object to expand into a neighbourhood, becoming a *bulk* rather than a sharp point. In this relaxed state, the object has more "energy" to interact with its environment: it collects information that may not strictly belong to the object itself but to the network it traverses. The neighbourhood functions as a probe of the medium, capturing how the object and the network resonate together.

In this sense, blurring is *catalytic*. It does not damage the object nor distort the network; instead, it enhances their interaction, allowing latent structure to surface. The messages that survive this joint evolution are then attached back to the object when the neighbourhood is removed (reading).

Thus, the philosophy of blur is:

> *Precision is achieved not by resisting relaxation, but by passing through it. Blurring reveals, interaction refines, and reading restores.*

In this sense, the applicability of blurring is *universal.*

Blur and Undecidability (toy but sharp)

We briefly show that blurring does *not* make hard decision problems easy: even a positive, normalized blur preserves the halting/non–halting gap.

Discrete time as a blurred object

Work in $\mathcal{C} = \mathbf{Meas}$. Let Prog be the set of programs (or initial states) and, for $e \in \mathsf{Prog}$, let $(F^t e)_{t\in\mathbb{N}}$ be its evolution with a halting predicate $h : \mathsf{Prog} \to \{0,1\}$ (true exactly on halting states). For $\lambda \in (0,1)$ define the time–blur kernel $K_\lambda(t) = (1-\lambda)\lambda^t$ on $\mathbb{N}$.

The *blurred halting mass* is the read-out

$$H_\lambda(e) \ := \ \sum_{t\geq 0} K_\lambda(t)\, h(F^t e) \ = \ \mathbb{E}_{T\sim\mathrm{Geom}(1-\lambda)}\big[h(F^T e)\big],$$

i.e. convolution in time with a positive, normalized kernel (a Markov/semigroup blur in the sense of our analytic model).

Theorem 13 (Undecidability survives blur)**.** *For every fixed $\lambda \in (0,1)$ and every program e with halting time $\tau(e) \in \mathbb{N} \cup \{\infty\}$,*

$$H_\lambda(e) = \begin{cases} \lambda^{\tau(e)} & \text{if } \ \tau(e) < \infty, \\ 0 & \text{if } \ \tau(e) = \infty. \end{cases} \qquad \textit{Hence} \quad \{\, e : \ H_\lambda(e) > 0 \,\} \ \equiv \ \mathsf{HALT}.$$

In particular, deciding the positivity *of this blurred read-out is undecidable.*

Sketch in our language. K_λ is a positive-definite mollifier on time (a *blur scale*). The evolution $e \mapsto (F^t e)_t$ is a blur-morphism into the path object; the observable h is a nonnegative map. If e halts at time τ, then $h(F^t e) = 0$ for $t < \tau$ and 1 thereafter, so $H_\lambda(e) = (1-\lambda)\sum_{t\geq\tau}\lambda^t = \lambda^\tau > 0$; if it never halts, the sum is 0. Thus the positivity read-out after blurring is many–one equivalent to HALT. □

Remark 14 (Fits the framework and links to computable analysis)**.** *This is a special case of our analytic blur model: B_λ is convolution on the time axis, $\bar{f}_\lambda$ is the (Markov) blur–morphism induced by the dynamics, and r^λ is the numerical read-out. The undecidability hinges only on* positivity, normalization*, and a strict positivity gap on halting—exactly the bloorf invariants. Conceptually this is close in spirit to classical encodings of halting into analytic properties of real-valued functions in computable analysis: blur does not eliminate undecidability; it packages it into a robust positivity gap.*

Continuous time and spatial blur (one-liners)

- **Continuous time.** With $K_y(t) = ye^{-yt}$ on $\mathbb{R}_{\geq 0}$ and halting time $\tau(e) \in [0, \infty]$,

$$H_y(e) = \int_0^\infty K_y(t)\, h(F^t e)\, dt = \mathbf{1}_{\{\tau(e)<\infty\}}\, e^{-y\tau(e)}.$$

  Again, $H_y(e) > 0$ iff e halts.

- **Spatial blur.** If halting emits a unit-mass bump at a known location x_e while non-halting emits 0, then any nonnegative kernel ϕ_y with $\phi_y(0) > 0$ yields $(\phi_y * f_e)(x_e) > 0$ iff e halts.

Moral. Blurring simplifies analysis by preserving positivity and behaving well with limits; those same features preserve the halting/non–halting dichotomy as a *strict* positivity gap. In the categorical language of blur, undecidability is therefore *bloorf–robust*: blur clarifies, it doesn't conjure answers.

Two guiding examples of blur

We briefly sketch two concrete examples that can be read directly in the language of blurred objects and blur–morphisms: (1) addition and multiplication as a blurred pair of channels; (2) Weyl equidistribution on the circle under Poisson blur. In both cases, the role of blur is not decorative: it is forced by the underlying harmonic analysis, and the categorical picture simply packages that necessity.

Addition and multiplication as a blurred pair

At a very classical level, addition and multiplication live on different backgrounds:

- addition acts naturally on $(\mathbb{R}, dy)$ via translations $T_a f(y) = f(y + a)$;
- multiplication acts naturally on $((0, \infty), dx/x)$ via dilations $D_c f(x) = f(cx)$.

A standard device is to move to the *log-line*: with $x = e^y$ the multiplicative group $(0, \infty)$ with Haar measure dx/x becomes $(\mathbb{R}, dy)$, and the Mellin transform on functions of x becomes the Fourier transform on functions of y. This is the precise sense in which multiplicative structure is a *blurred* copy of additive structure.

Fix the base category $\mathcal{C} = \mathbf{Meas}$ (or, if preferred, the corresponding L^2 Hilbert spaces). Consider two blurred objects:

$$\mathbf{X}_\times := ((0, \infty), \mathcal{B}^\times), \qquad \mathbf{X}_+ := (\mathbb{R}, \mathcal{B}^+),$$

where:

- $B^{\times}_{\varepsilon}$ is *log–Gaussian blur* on the multiplicative line: if $U : L^2((0,\infty), dx/x) \to L^2(\mathbb{R}, dy)$ is the log-change isometry $(Uf)(y) = e^{y/2} f(e^y)$, and G_ε is the Gaussian of variance ε, then

$$B^{\times}_{\varepsilon} f := U^{-1}\big(G_\varepsilon * (Uf)\big),$$

  i.e. blur in log-coordinates, pulled back to the x-line;

- B^{+}_{τ} is ordinary Gaussian blur on the additive line:

$$B^{+}_{\tau} g := G_\tau * g, \qquad g : \mathbb{R} \to \mathbb{C}.$$

The soft addition / soft multiplication picture can then be read as follows.

- On the additive side, $+$ is sharp: we are free to work with $g_1 \oplus g_2$ defined by $(g_1 \oplus g_2)(y) = g_1(y) + g_2(y)$ or with convolution in y, depending on the level of structure we track.
- On the multiplicative side, $\times$ is sharp: we work with pointwise product in x or with multiplicative convolution

$$(f_1 *_\times f_2)(x) = \int_0^\infty f_1\Big(\frac{x}{u}\Big) f_2(u)\, \frac{du}{u}.$$

- The link between the two channels is the isometry U, which intertwines multiplication on $(0,\infty)$ with translation in the spectral variable on the log-line.

The crucial fact is that Fourier on the log-line satisfies a genuine Heisenberg-type uncertainty: if $\widehat{h}(\xi)$ is the Fourier transform of $h(y)$, then one cannot simultaneously have both h and $\widehat{h}$ sharply localized. Concretely, there is a lower bound of the form

$$\mathrm{Var}_y(h)\, \mathrm{Var}_\xi(\widehat{h}) \;\geq\; \frac{1}{4}$$

(up to the usual normalization choices). On the multiplicative side this reads: you cannot make the blur in log-x arbitrarily small while also keeping the multiplicative spectral content arbitrarily sharp.

In the language of blurred objects:

- the families $(B^{\times}_{\varepsilon})_{\varepsilon>0}$ and $(B^{+}_{\tau})_{\tau>0}$ are blur functors on $\mathbf{X}_\times$ and $\mathbf{X}_+$;
- the log-change U induces a blur-morphism

$$\overline{\mathrm{Switch}}_{\times\to+} :\ \mathbf{X}_\times \longrightarrow \mathbf{X}_+, \qquad \big(\overline{\mathrm{Switch}}_{\times\to+}\big)_\varepsilon := B^{+}_{\tau(\varepsilon)} \circ U \circ B^{\times}_{\varepsilon},$$

  where the blur indices are coupled so that $\varepsilon \cdot \tau(\varepsilon) \geq c > 0$ for a universal constant.

The Heisenberg inequality is exactly the statement that we cannot make *both* ε and $\tau(\varepsilon)$ arbitrarily small: there is a hard lower bound on the blur budget required to switch between the additive and multiplicative lenses. In other words, any blur-morphism that genuinely intertwines the two channels must pay with a nonzero uncertainty window.

Soft addition and soft multiplication then appear as operations defined *at the blurred level*:

- to add multiplicative data "as if" it were additive, we
  1. blur in the multiplicative channel at some ε,
  2. switch to the additive channel via $\overline{\mathrm{Switch}}_{\times\to+}$,
  3. perform the native additive operation there,

4. and read back sharply.

- to multiply additive data "as if" it were multiplicative, we do the symmetric construction using the inverse blur-morphism $\overline{\mathtt{Switch}}_{+\to\times}$.

The *Grand Lemma* then says precisely: sending sharp data through this network is equivalent to sending it together with its neighbourhood and reading in successive steps. The analytic content (Fourier–Mellin duality and Heisenberg uncertainty) ensures that such a network cannot exist with zero blur; the categorical content ensures that, once we fix a blur budget, properties that are stable under retracts and filtered limits can be transported between the two channels.

Thus the addition/multiplication pair is a particularly vivid example of a blurred object and blur–morphism that are not optional: blur is not something we choose to add; it is something the harmonic analysis forces us to account for.

Weyl equidistribution under Poisson blur

As a second example, we take a classical theorem where blur appears almost unavoidably in standard proofs: Weyl's equidistribution of $\{n\alpha\}$ modulo 1.

Let $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$ be the circle with Lebesgue probability measure m, and fix $\alpha \in \mathbb{R}$. Consider the orbit $x_n := n\alpha \pmod{2\pi}$ and the empirical measures

$$\mu_N := \frac{1}{N}\sum_{n=1}^{N} \delta_{x_n}.$$

Weyl's theorem says: if $\alpha/2\pi$ is irrational, then μ_N converges weak-* to m, i.e. the orbit is equidistributed.

In our language, we define a blur on $\mathbb{T}$ via the Poisson kernel

$$P_r(\theta) \;=\; \frac{1-r^2}{1-2r\cos\theta + r^2}, \qquad 0<r<1,$$

and set $B_r f := P_r * f$ for $f \in L^1(\mathbb{T})$. The family $(B_r)_{0<r<1}$ is a standard approximate identity:

$$\lim_{r\uparrow 1} B_r f = f \quad \text{in } L^p(\mathbb{T})$$

for $1 \le p < \infty$, and B_r annihilates high Fourier modes exponentially: if $f(\theta) = \sum_{k\in\mathbb{Z}} \hat{f}(k)e^{ik\theta}$, then

$$B_r f(\theta) = \sum_{k\in\mathbb{Z}} r^{|k|}\hat{f}(k)e^{ik\theta}.$$

Define blurred empirical measures

$$\mu_{N,r} := \mu_N * P_r, \qquad 0<r<1,$$

so that for a continuous f,

$$\int_{\mathbb{T}} f\, d\mu_{N,r} = \int_{\mathbb{T}} (B_r f)\, d\mu_N = \frac{1}{N}\sum_{n=1}^{N} (B_r f)(x_n).$$

Now set up the blurred object

$$\mathbf{X}_{\mathbb{T}} := (\mathbb{T}, \mathcal{B}), \qquad B_r := \mathcal{B}(r) :\; L^1(\mathbb{T}) \to L^1(\mathbb{T}),$$

with $r \in I := (0,1)$, partially ordered by $r \preceq r'$ if $r \le r'$. The sharp limit ρ is $f \mapsto \lim_{r\uparrow 1} B_r f = f$ (when it exists), and the exhaustion at "infinity" is the collapse to constants as $r \downarrow 0$.

The property of interest is:

$$\mathsf{P}(\mu_\bullet): \quad \mu_N \xrightarrow{w^*} m \text{ as } N \to \infty \quad \text{(equidistribution)}.$$

Its blurred version at scale r is:

$$\mathsf{P}_r(\mu_\bullet): \quad \mu_{N,r} \xrightarrow{w^*} m \text{ as } N \to \infty,$$

i.e. equidistribution holds after Poisson blur at fixed $r < 1$.

Classically, Weyl's criterion reduces $\mathsf{P}(\mu_\bullet)$ to the decay of exponential sums:

$$\frac{1}{N}\sum_{n=1}^{N} e^{ikx_n} \longrightarrow 0 \quad (N \to \infty) \qquad \text{for all } k \in \mathbb{Z} \setminus \{0\}.$$

In the blurred picture, we test against $B_r f$ instead. On Fourier coefficients, this inserts the extra factor $r^{|k|}$, making the estimates strictly easier:

$$\frac{1}{N}\sum_{n=1}^{N}(B_r f)(x_n) = \sum_{k\in\mathbb{Z}} r^{|k|}\hat{f}(k)\Big(\frac{1}{N}\sum_{n=1}^{N} e^{ikx_n}\Big).$$

For each fixed $r < 1$, the tail over large $|k|$ is exponentially suppressed, and the analysis reduces to finitely many frequencies with a built-in damping factor $r^{|k|}$. In other words, $\mathsf{P}_r(\mu_\bullet)$ is strictly easier to verify than $\mathsf{P}(\mu_\bullet)$.

From the viewpoint of the property-transport Proposition (transport under uncertainty window), equidistribution is:

- stable under convolution with an approximate identity (Poisson blur): if $\mu_{N,r} \to m$ for a fixed $r < 1$ and $B_r \to \mathrm{id}$ as $r \uparrow 1$, then $\mu_N \to m$;
- stable under filtered limits in r and monotone with respect to the natural preorder on kernels ($P_r \precsim P_{r'}$ if $r \le r'$ in the sense of pointwise domination).

Thus, in this example, we may legitimately work at any fixed blur level $r < 1$, prove $\mathsf{P}_r(\mu_\bullet)$ in the blurred world (where the Fourier side is better behaved), and then let $r \uparrow 1$ to recover the sharp statement $\mathsf{P}(\mu_\bullet)$. The Poisson blur is not an arbitrary smoothing trick: it is precisely a blur functor B_r with a reading map ρ that satisfies the hypotheses of our transport principle.

In summary:

- addition vs. multiplication illustrates how blur is forced by Fourier–Mellin uncertainty when we try to switch between algebraic channels;
- Weyl equidistribution on the circle illustrates how blur is a natural technical lens: we prove a property in a blurred regime where the analysis is easier, and then read back to the sharp regime using stability under limits.

Both examples fit cleanly into the categorical framework developed above and show that blur is not merely a philosophical decoration, but a mathematically inevitable mediator in situations where different structures must be made to communicate.

Blur at closure interfaces: from successors to $\mathbb{N}$ and from $\mathbb{Q}$ to $\mathbb{R}$

We briefly record two concrete interfaces where blur is not an optional decoration but the *mechanism* by which an idealized infinite object is constructed from a finitary one:

- the passage from a single successor step "+1" to the full set of natural numbers $\mathbb{N}$;
- the passage from the rationals $\mathbb{Q}$ to the reals $\mathbb{R}$ via Cauchy completion.

Both can be phrased as instances of our general notion of a blurred object $\mathbf{X} = (X, \mathcal{B}_X)$ with a faithful reading cone $\rho_X : \lim_I \mathcal{B}_X \xrightarrow{\cong} X$.

Successor orbits and the finite–infinite interface

Fix a base category $\mathcal{C} = \mathbf{Set}$ (or $\mathbf{Meas}$ with counting measure). A *successor orbit* is a triple (O, S, o_0) where O is a set, $o_0 \in O$ is a distinguished point, and $S : O \to O$ is an injective map such that every element lies on the forward orbit

$$O \;=\; \{\, S^n(o_0) : n \in \mathbb{N} \,\},$$

and there are no cycles (i.e. $S^n(o_0) = S^m(o_0)$ implies $n = m$).

In [15] we considered the *Abel blur* on such an orbit: for each $q \in (0, 1)$ define a probability measure

$$\mu_q \;:=\; (1-q) \sum_{n \geq 0} q^n \, \delta_{S^n(o_0)},$$

and regard $\mathcal{B}_O : (0, 1) \to \mathcal{C}$ as the functor sending q to the blurred object $B_q O$ that carries these weights. Reading at scale q is integration against μ_q, and the sharp reading ρ_O is the Abel-limit functional as $q \uparrow 1$ (when it exists). In this language $\mathbf{O}$ is a blurred object with index set $I = (0, 1)$ and reading cone $\rho_O : \lim_{q \uparrow 1} B_q O \to O$.

The main result there can be summarized as follows.

Theorem 15 (Successor orbit $\iff$ Abel blur on time)**.** *Let (O, S, o_0) be a successor orbit without cycles. Then the following are equivalent:*

1. *There is a bijection $\varphi : O \to \mathbb{N}$ such that $\varphi(Sx) = \varphi(x) + 1$ and $\varphi(o_0) = 0$.*
2. *O carries a faithful Abel blur $(B_q O)_{q \in (0,1)}$ with:*
   - *(a) each $B_q O$ supported on $\{S^n(o_0)\}_{n \geq 0}$ with weights $(1-q)q^n$;*
   - *(b) S acts as a blur-morphism $\overline{S} : \mathbf{O} \to \mathbf{O}$ commuting with the blur scales (shift of time);*
   - *(c) the associated reading cone ρ_O is faithful, $O \simeq \lim_{q \uparrow 1} B_q O$.*

In other words, having a copy of $\mathbb{N}$ is equivalent to having a canonical Abel blur on the successor orbit that survives all blur scales and recovers the orbit in the sharp limit.

Remark 16 (Finite–infinite interface as blur)**.** *The measures μ_q are exactly the* resolvent blur *of the discrete-time dynamics S—a positive, normalized kernel on time (a geometric law). The passage from a single successor step to the full $\mathbb{N}$ is mediated by this blur: the orbit is probed by ever-longer geometric windows as $q \uparrow 1$, and the sharp object is recovered as the limit of these blurred views. No extra structure beyond blur is needed to pass from "one more step" to an infinite chain; the Abel blur is the interface.*

Cauchy blur and the completion of $\mathbb{Q}$

We now look at the familiar closure $\mathbb{Q} \subset \mathbb{R}$ and show that it is equally natural to regard $\mathbb{R}$ as the sharp reading of a blurred object constructed from $\mathbb{Q}$. For definiteness, take $\mathcal{C} = \mathbf{Meas}$ (or $\mathbf{Met}$) with the usual absolute-value metric.

Fix a scale index set $I = (0, \infty)$ with order $\varepsilon' \preceq \varepsilon$ iff $\varepsilon' \leq \varepsilon$ and monoidal sum given by $\varepsilon \oplus \delta = \varepsilon + \delta$. For each $\varepsilon > 0$ define $B_\varepsilon \mathbb{Q}$ to be the set of ε-*Cauchy profiles*: families $(q_\delta)_{\delta \in (0,\varepsilon]}$ with $q_\delta \in \mathbb{Q}$ and

$$|q_\delta - q_{\delta'}| \;\leq\; \max\{\delta, \delta'\} \qquad \text{for all } \delta, \delta' \in (0, \varepsilon].$$

If $\varepsilon' \leq \varepsilon$, the structure map $B_\varepsilon \mathbb{Q} \to B_{\varepsilon'} \mathbb{Q}$ is the restriction $(q_\delta)_{\delta \leq \varepsilon} \mapsto (q_\delta)_{\delta \leq \varepsilon'}$.

Uncertainty window. We equip $B_\varepsilon \mathbb{Q}$ with the canonical thickening and read-off maps

$$j^\varepsilon_{\mathbb{Q}} : \; \mathbb{Q} \to B_\varepsilon \mathbb{Q}, \qquad j^\varepsilon_{\mathbb{Q}}(q) := (q_\delta \equiv q)_{\delta \in (0,\varepsilon]},$$

and

$$r^{\varepsilon}_{\mathbb{Q}} : \; B_\varepsilon\mathbb{Q} \to \mathbb{Q}, \qquad r^{\varepsilon}_{\mathbb{Q}}\big((q_\delta)_\delta\big) := q_\varepsilon$$

(read-off at the outermost scale). Clearly $r^{\varepsilon}_{\mathbb{Q}} \circ j^{\varepsilon}_{\mathbb{Q}} = \mathrm{id}_{\mathbb{Q}}$. Moreover, $(j^{\varepsilon}_{\mathbb{Q}}, r^{\varepsilon}_{\mathbb{Q}})$ is a lax retract in the sense of Definition 4 once we endow $B_\varepsilon\mathbb{Q}$ with the natural "fuzz" preorder

$$(q_\delta)_\delta \;\precsim\; (p_\delta)_\delta \quad :\Longleftrightarrow \quad \sup_{\delta\in(0,\varepsilon]} |q_\delta - p_\delta| \;\le\; \varepsilon.$$

Indeed, for any profile $(q_\delta)_\delta$ and any $\delta \le \varepsilon$, the defining Cauchy condition gives $|q_\delta - q_\varepsilon| \le \varepsilon$, hence

$$\sup_{\delta\in(0,\varepsilon]} |q_\delta - q_\varepsilon| \;\le\; \varepsilon,$$

so the constant profile at q_ε lies ε-below the original:

$$j^{\varepsilon}_{\mathbb{Q}}\Big(r^{\varepsilon}_{\mathbb{Q}}((q_\delta)_\delta)\Big) \;=\; (q_\varepsilon)_\delta \;\precsim\; (q_\delta)_\delta.$$

This expresses the intended meaning: "read at scale ε, then thicken back" cannot recover micro-variation inside the ε-cell.

We thus form the blurred object

$$\mathbf{Q}_{\text{Cauchy}} := (\mathbb{Q}, \mathcal{B}_{\mathbb{Q}}), \qquad \mathcal{B}_{\mathbb{Q}}(\varepsilon) := B_\varepsilon\mathbb{Q}.$$

Proposition 17 (Sharp reading gives the completion)**.** *Let X be a complete metric space containing $\mathbb{Q}$ as a dense subspace via an isometric embedding $i : \mathbb{Q} \hookrightarrow X$. Then:*

1. *There is a canonical reading map*

$$\rho_X : \lim_{\varepsilon\downarrow 0} B_\varepsilon\mathbb{Q} \to X$$

sending each compatible Cauchy profile to its limit in X.

2. *ρ_X is an isometry and is surjective. It becomes an isometric isomorphism after identifying profiles that are 0-distance apart, i.e. after passing to* blur-classes*:*

$$X \;\simeq\; \Big(\lim_{\varepsilon} B_\varepsilon\mathbb{Q}\Big)/\!\sim, \qquad (q_\varepsilon)_\varepsilon \sim (p_\varepsilon)_\varepsilon \iff \lim_{\varepsilon\downarrow 0} |q_\varepsilon - p_\varepsilon| = 0.$$

In particular, the completion X of $\mathbb{Q}$ is precisely the sharp reading *of the Cauchy blur diagram $(B_\varepsilon\mathbb{Q})_{\varepsilon>0}$, once we forget the inessential "micro-history" of choosing representatives.*

Sketch. Compatibility under restriction means that an element of $\lim_\varepsilon B_\varepsilon\mathbb{Q}$ is equivalently a single family $(q_\varepsilon)_{\varepsilon>0}$ with

$$|q_\varepsilon - q_{\varepsilon'}| \;\le\; \max\{\varepsilon, \varepsilon'\} \qquad (\varepsilon, \varepsilon' > 0),$$

i.e. a Cauchy approximation scheme with a built-in modulus. Since $i(\mathbb{Q})$ is dense and X is complete, $(i(q_\varepsilon))_\varepsilon$ converges to a unique point $x \in X$; define $\rho_X((q_\varepsilon)_\varepsilon) := x$. If two such families satisfy $|q_\varepsilon - p_\varepsilon| \to 0$, then $i(q_\varepsilon) - i(p_\varepsilon) \to 0$, so they have the same limit; thus ρ_X factors through the quotient by $\sim$. Conversely, any $x \in X$ admits a rational approximation (q_ε) with $|i(q_\varepsilon) - x| \le \varepsilon$; one checks this produces an element of the limit, and ρ_X sends it to x. Standard uniqueness of metric completions then yields the claimed isometric identification. $\square$

Remark 18 (Blur is the mechanism behind completion)**.** *From this viewpoint, completion is exactly the slogan:* to read $\mathbb{Q}$ sharply, one must first blur it coherently across scales. *A real number is not a single rational, but a compatible family of ε-cells; different microscopic choices of representatives define the same sharp point once the blur goes to 0. In particular, the familiar passage $\mathbb{Q} \subset \mathbb{R}$ is a sharp reading of a blurred object: the "loss of micro-history" is not a defect but the categorical content of completion.*

Closure as controlled information loss

Both interfaces—successor $\to \mathbb{N}$ and $\mathbb{Q} \to \mathbb{R}$—exhibit the same pattern:

- At the *finitary* level we have a concrete mechanism: a successor map S or a dense countable set $\mathbb{Q}$.
- At the *idealized* level we have an infinite object: the chain $\mathbb{N}$ or the complete line $\mathbb{R}$.
- The passage from the former to the latter is always mediated by a *blur diagram*: Abel blur on time or Cauchy blur on rationals.

In both cases, some information is irrevocably discarded:

- in the successor case, the detailed placement of finitely many steps is invisible to the Abel blur as $q \uparrow 1$; all orbits that are *eventually the same* produce the same sharp object;
- in the Cauchy case, the detailed rational path taken to approach a real number is invisible to the completion: all Cauchy families that converge to the same $x \in \mathbb{R}$ are identified.

The blur calculus precisely isolates this loss: it lives entirely in the neighbourhood part of the blurred object $(X, \mathcal{B}_X)$, while the sharp object X itself remains clean. The point is not to undo the information loss, but to *track where it happens and with what budget.*

Proposition 19 (Blur inevitability for metric completions)**.** *Let (X, d) be a metric space and $\widehat{X}$ a metric completion. Then there exists a blurred object $\mathbf{X} = (X, \mathcal{B}_X)$ of Cauchy type such that:*

1. *$\widehat{X} \simeq \lim_{\varepsilon \downarrow 0} B_\varepsilon X$ as in Proposition 17;*
2. *any other completion of X arises from an isomorphic blur diagram.*

Hence every metric completion can be realized, up to unique isometry, as a sharp reading of a canonical blur on X.

Sketch. Replace $\mathbb{Q}$ by X in the construction of Subsection : $B_\varepsilon X$ is the set of ε-Cauchy profiles in X, with restriction maps as above. Any completion $\widehat{X}$ of X is, by definition, obtained by quotienting the Cauchy families by the indistinguishability relation "converge to the same limit". This is exactly the identification performed by the limit $\lim_\varepsilon B_\varepsilon X$ and the reading cone $\rho_{\widehat{X}}$. Uniqueness of completions up to isometry yields the second statement. ☐

Remark 20 (Which idealizations hide blur?)**.** *Idealizations that present themselves as* closures *or* completions *—metric completion, topological closure via nets/filters, measure-theoretic completion modulo null sets—all share the same pattern: points of the ideal object are equivalence classes of blurred histories in the base object. The categorical blur package makes this explicit. In particular, whenever an idealization kills micro-history by identifying different approximating families, blur is not optional: it is the natural language for what has been lost and why the resulting sharp object is still well-behaved.*

Together with the earlier analytic and logical models, these two closure interfaces show that *blur is not a taste in notation.* It sits exactly at the junctions where mathematics passes from finite to infinite, from discrete to continuous, or from partial information to idealized objects. The category $\mathbf{Blur}(\mathcal{C})$ is therefore not an arbitrary enrichment, but a minimal setting in which such passages can be tracked, controlled, and—when possible—reversed up to the natural blur equivalences.

References

[1] F. W. Lawvere, Metric spaces, generalized logic, and closed categories, *Rendiconti del Seminario Matematico e Fisico di Milano* **43** (1973), 135–166.

[2] G. M. Kelly, *Basic Concepts of Enriched Category Theory*, London Mathematical Society Lecture Note Series, vol. 64, Cambridge University Press, 1982.

[3] S. Abramsky and A. Jung, Domain theory, in S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum (eds.), *Handbook of Logic in Computer Science*, vol. 3, Oxford University Press, 1994, pp. 1–168.

[4] C. Jones and G. D. Plotkin, A probabilistic powerdomain of evaluations, in *Proceedings of the 4th IEEE Symposium on Logic in Computer Science (LICS)*, 1989, pp. 186–195.

[5] M. Giry, A categorical approach to probability theory, in B. Banaschewski (ed.), *Categorical Aspects of Topology and Analysis*, Lecture Notes in Mathematics, vol. 915, Springer, 1982, pp. 68–85.

[6] T. Fritz, A synthetic approach to Markov kernels, conditional independence, and theorems on sufficient statistics, *Advances in Mathematics* **370** (2020), 107239.

[7] K. Weihrauch, *Computable Analysis*, Springer, 2000.

[8] M. B. Pour-El and J. I. Richards, *Computability in Analysis and Physics*, Springer, 1989.

[9] A. M. Turing, On computable numbers, with an application to the Entscheidungsproblem, *Proceedings of the London Mathematical Society* (2) **42** (1936), 230–265.

[10] H. Weyl, Über die Gleichverteilung von Zahlen mod. Eins, *Mathematische Annalen* **77** (1916), 313–352.

[11] W. Rudin, *Real and Complex Analysis*, 3rd ed., McGraw–Hill, 1987.

[12] E. M. Stein and R. Shakarchi, *Fourier Analysis: An Introduction*, Princeton University Press, 2003.

[13] A. Perišić, *Soft-Addition and Soft-Multiplication and the Channel–Switch Error*, Zenodo, 2025.

[14] A. Perišić, *Blur as a Universal Principle*, Zenodo, 2025.

[15] A. Perišić, *Blur at the Finite–Infinite Interface*, Zenodo, 2025.

[16] A. Perišić, *The Reals as Blur–Completion of the Rationals*, Zenodo, 2025.