

Text S1. Construction of Fig.5 Data for the bubble plot were extracted from the studies included in this review, with four variables recorded per system: catalyst dose (g L^{-1}), malachite green (MG) removal efficiency (%), reaction time (min), and irradiation source. Where a study reported efficiency under more than one condition, material form, or analogue, the highest value stated by the authors was used (e.g. ZnVFeO_4 , 98% at 25 ppm rather than 92% at 100 ppm; Ti-doped BiVO_4 , 100% for the powder rather than 96% for the polyurethane-supported film; Cu-MOFs, 98% for Cu(Seba) rather than 97% for Cu(Azla)). Figure 5 therefore represents best reported performance, not performance under a standardised test condition. Doses given with units embedded in the text (e.g. "0.3 g/L") were converted to numeric values, and times reported in hours were converted to minutes. One system (MG-MI, Bi_2O_3 -doped TiO_2 on molecularly imprinted powdered activated carbon) was excluded because no catalyst dose was reported, leaving 23 entries.

Irradiation sources were aggregated into four colour categories for legibility, since the seven source types distinguished in Fig. 3 would not resolve visually at this scale: **UV** (UV lamps and the 365 nm UV LED), **visible** (Hg, Xe, tungsten, and halogen lamps operating in the visible range), **natural sunlight** (direct solar irradiation), and **combined UV–visible** (studies reporting performance under both). The unaggregated source for every entry is retained in Online Resource 1.

The plot was generated in Microsoft Excel: catalyst dose on the x-axis (logarithmic scale, to resolve the cluster of systems dosed below 0.5 g L^{-1}), MG removal efficiency on the y-axis, bubble area proportional to reaction time, and bubble colour denoting irradiation category. Each catalyst was plotted as a discrete series so that individual systems could be labelled directly. The underlying dataset is provided in Online Resource 1.

Radar plot:

Multi-metric comparison was performed by rescaling six performance indicators to a common 0–1 range so that catalysts could be compared on a single radar axis system despite differences in reported units. Removal efficiency was expressed on its natural bounded scale (fraction of 1) rather than min–max normalized, since percentage removal is already dimensionless and min–max rescaling of a saturated variable would exaggerate trivial differences. Apparent rate constant (k , min^{-1}) and adsorption capacity (mg g^{-1}) were min–max normalized across the 24 included studies, so that 1.000 denotes the highest reported value and 0.000 the lowest. Reaction time was inverted before normalization, $(1 - [(t - t_{\text{min}})/(t_{\text{max}} - t_{\text{min}})])$, so that 1.000 denotes the shortest reaction time and higher scores consistently indicate better performance across all axes. Reusability was normalized against the maximum number of cycles reported in the dataset (11 cycles). Irradiation source was assigned an ordinal score reflecting its suitability for scalable deployment. Where a metric was not reported, the cell was left empty rather than imputed as zero, so the corresponding radar vertex shows a gap; reuse data were unavailable for three catalysts

(CTAS@NiFe₂O₄, NFO@C, and MG-MI). All normalization was performed on the full 24-study dataset, and the underlying values and formulas are provided in Online Resource 1.

Text S2. Normalization and construction of Fig. 6

Multi-metric comparison was performed by rescaling six performance indicators to a common 0–1 range so that catalysts could be compared on a single radar axis system despite differences in reported units. Removal efficiency was expressed on its natural bounded scale (fraction of 1) rather than min–max normalized, since percentage removal is already dimensionless and min–max rescaling of a saturated variable would exaggerate trivial differences. Apparent rate constant (k , min⁻¹) and adsorption capacity (mg g⁻¹) were min–max normalized across the 24 included studies, so that 1.000 denotes the highest reported value and 0.000 the lowest. Where k was not reported by the original authors, it was recalculated assuming pseudo-first-order kinetics with conversion capped at 0.9999 (see §4.6); such values are conservative approximations rather than measured constants and are flagged in Online Resource 1. Reaction time was inverted before normalization, $(1 - [(t - t_{\min})/(t_{\max} - t_{\min})])$, so that 1.000 denotes the shortest reaction time and higher scores consistently indicate better performance across all axes. Reusability was normalized against the maximum number of cycles reported in the dataset (11 cycles; Tanuj et al. 2024). Irradiation source was assigned an ordinal score reflecting its suitability for scalable deployment.

Where a metric was not reported, the cell was left empty rather than imputed as zero, so the corresponding radar vertex shows a gap. Reuse cycle counts were unavailable for three catalysts: NFO@C, for which reuse was not tested, and CTAS@NiFe₂O₄ and MG-MI, for which the authors reported qualitative stability after reuse without stating the number of cycles. These were deliberately not imputed as zero, since a gap accurately represents absent data whereas zero would imply measured failure. All normalization was performed on the full 24-study dataset, and the underlying values and formulas are provided in Online Resource 1.

The radar plot was constructed in Excel by plotting column A (Catalyst) against six data columns: I, J, K, L, M, and O. To reproduce it, select A1:A25, then Ctrl-select I1:M25 and O1:O25, and insert a Radar chart of the "Radar with markers" type. The six resulting axes correspond to those listed in the Fig. 6 caption: efficiency (column I, Efficiency_abs), kinetics (column J, k_{norm}), adsorption capacity (column K, Cap_norm), reaction time (column L, Time_score, inverted so that 1.000 denotes the fastest reaction), reusability (column M, Reuse_norm), and light source (column O, Lamp_score). For the four min–max normalized axes (J, K, L, and M), the radar rings at 0.2, 0.4, 0.6, and 0.8 correspond to 20%, 40%, 60%, and 80% of the observed range. Column I is plotted on its own natural 0–1 scale rather than being min–max normalized; this is documented in the cell comment on I1 and stated in the figure caption.