© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Graph Neural Network Prediction of Antioxidant Activity of Polyphenols from Grape Pomace

ADITI EKHANDE, NOVYTE MATERIALS

Abstract*—Grape pomace is the main solid by-product of winemaking and a cheap source of polyphenols. However, the antioxidant strength of most compounds in grape pomace has never been measured on its own. Mixtures do not behave additively, so extract results cannot be separated into compound-level values. This paper reviews the chemistry and the modelling problem, sets out a full pipeline from data curation to molecular graphs to a trained network, and then runs it on measured data. Three results concern the data. Unit conversion matters, since epicatechin measured in two laboratories in two-unit systems agrees to 0.021 pIC_{50} log units once converted. Donor count does not order activity, since gallic acid and trans-resveratrol carry the same three phenolic hydroxyls yet differ by 0.436 log units, which is the argument for a representation that encodes adjacency. And protocols cannot be pooled, since three compounds measured in a second laboratory differ systematically by 1.152 log units, more than the 1.477 log unit range of the training corpus. For training, 25 phenolics measured in one laboratory under one protocol reduced to 16 usable records once non-numeric entries were removed and dose-response fit quality was filtered on, a filter that published curation protocols do not apply and that removed a compound whose tabulated IC_{50} came from a regression explaining 19 percent of its variance. On those 16 compounds a message-passing network reached a leave-one-out R^2 of 0.744 against 0.138 for a mean predictor, with a training fit of 0.696, and a label-scrambling control confirmed that no structure-activity signal was recovered at this sample size. Predictions for the five targets span 0.322 log units and all ten pairwise comparisons overlap at 95 percent. The pipeline runs correctly, and closing the gap between 16 compounds and the 1911 used by published benchmarks is the work that remains.

***Keywords:* Grape pomace, Polyphenols, Antioxidant activity, DPPH**

Graphical abstract

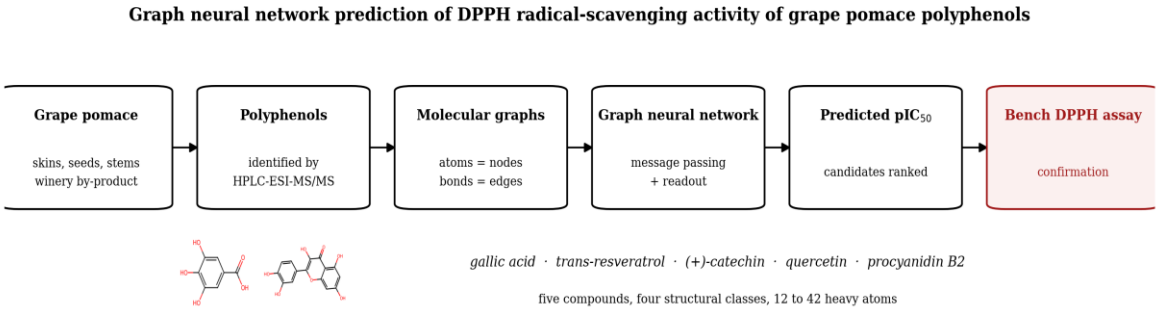

I. INTRODUCTION

Grape pomace, or marc, is what remains after grapes are pressed for wine or juice. It is made up of skins, seeds and, depending on the process, stems, and it accounts for a large share of the fruit mass. Production is seasonal and concentrated at individual wineries, and most of it is composted, spread on land or sent to animal feed. Because the residue is wet, acidic and rich in polyphenols, effluent from it carries a high oxygen demand, so disposal is a cost in its own right [1]. Pressing removes the sugar and juice but leaves most of the phenolic material behind, which makes the pomace a low-cost feedstock for compounds that the food, cosmetic and nutraceutical industries already buy. Recovering them is a recurring objective in the circular-economy literature [2, 3]. The phenolic content of pomace falls into a small number of structural classes (Table 2). Phenolic acids divide into hydroxybenzoic types such as gallic, protocatechuic and syringic acid and hydroxycinnamic types such as caffeic, p-coumaric and ferulic acid. Flavan-3-ols include the monomers (+)-catechin and (-)-epicatechin together with the procyanidin dimers B1 to B4 and larger condensed tannins. Flavanols are dominated by quercetin and its glycosides, with kaempferol and myricetin derivatives also present. Anthocyanins occur in red

cultivars as the 3-O-glucosides and acylated glucosides of malvidin, peonidin, delphinidin, petunidin and cyanidin. Stilbenes are represented by trans-resveratrol, its glucoside piceid, and piceatannol.

Two quantitative patterns recur. Catechin is usually the most abundant non-anthocyanin compound: it made up 49 to 73% of the monomeric flavan-3-ol fraction across four white cultivars [4] and reached 150 mg per 100 g in Brazilian red pomace [5]. In red pomace the anthocyanins are the largest fraction by mass, at 497 to 2918 mg per 100 g depending on cultivar [5]. Composition varies with grape variety, growing conditions, the winemaking route, the drying method and the extraction step. Seeds and skins differ systematically, with seeds carrying the flavan-3-ols and skins the anthocyanins and flavonols, so even the ratio of the two fractions in a batch alters the profile [4, 6]. An extract contains dozens of phenolics at once, and an assay on the whole extract cannot say which of them is responsible. Answering that needs a purified standard and a separate assay for every compound, which is slow and expensive, and only possible for compounds that are sold commercially. Arithmetic on extract data cannot close the gap either. Tournour et al. [1] mixed three cultivars and found the mixture gave lower total phenolic content and lower ORAC than the values predicted from the individual cultivars, and Pinelo et al. [7] reported the same negative interaction for mixtures of pure phenolics. Zhu et al. [3] made a related point: grape pomace had higher total phenolic content than kiwi pomace yet weaker hydroxyl-radical scavenging, so activity depends on which phenolics are present and not on the total alone. Compound-level values therefore have to be measured or predicted.

Antioxidant properties of grape pomace have been studied extensively, and machine learning has been applied to molecular activity prediction for years, but the two lines of work have rarely been joined. Molecular graph models have reached antioxidant endpoints mostly through peptides and mostly as classification, rather than through plant polyphenols as regression. Most published antioxidant models predict from precomputed descriptors, so the link runs from structure to descriptor to activity instead of from structure to activity. Graph models on this endpoint have not been benchmarked against conventional models on identical splits, so it is not known whether the added complexity pays for itself. And no published model has been used to prioritise the specific compounds that chromatography actually finds in pomace. Sections 2 to 6 set out what a study closing these gaps needs to do. They compile and curate measured antioxidant data for structurally diverse compounds, build molecular graphs from those structures, specify a graph neural network for DPPH activity, compare it against conventional models, apply the result to selected pomace polyphenols, and identify which structural features the predictions track. Section 7 then carries the pipeline out on a small curated corpus and reports what it produced, together with the controls that show how far the result can be pushed.

Table 1
Extraction methods applied to grape pomace and their reported behavior. Entries describe tendencies. There is no ranking, because the best condition depends on compound class and on the endpoint measured.

Method	Principle	Reported strength	Main caution
Solid-liquid (SLE)	Solvent contact with milled solid, often aqueous ethanol	Cheap, general, most widely reported	Long hot contact degrades phenolics; yield sensitive to solvent, time and ratio
Ultrasound-assisted (UAE)	Cavitation collapses bubbles against the cell wall	Short times; often the richest list of resolved phenolics	Local heating and radical formation; probe geometry affects reproducibility
Microwave-assisted (MAE)	Dielectric heating of solvent inside the matrix	Rapid; effective for tannins and anthocyanins	Anthocyanins degrade if power and duration uncontrolled
Supercritical fluid (SFE)	Supercritical CO_2 with a polar co-solvent	Solvent-free product; good for total polyphenols and flavonoids	Neat CO_2 poorly suited to polar polyphenols; high capital cost

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Method	Principle	Reported strength	Main caution
Pressurized water	Raised temperature and pressure lower water polarity	Green and solvent-free	Thermal degradation at the temperatures needed
Enzyme-assisted	Cell-wall enzymes release matrix-bound phenolics	Recovers bound phenolics solvent extraction leaves behind	Enzyme cost and specificity; extra process variables
Deep eutectic solvents	Designer solvents tuned to polyphenol polarity	High efficiency, low toxicity	Downstream solvent removal; poor comparability with older work

Moutinho et al. [2] optimized ethanol concentration for one pomace against six endpoints and obtained six different optima, from 53.0% for DPPH to 90% for ortho-diphenols, which is why no single setting can be called best.

Table 2
Phenolic classes reported in grape pomace, with the structural feature relevant to radical scavenging.

Class	Representative compounds	Main tissue	Feature relevant to scavenging	Graph size (heavy atoms)
Hydroxybenzoic acids	Gallic, protocatechuic, syringic, vanillic acid	Seed and skin	Two or three adjacent hydroxyls on one ring plus a carboxyl	11–15
Hydroxycinnamic acids	Caffeic, p-coumaric, ferulic acid	Skin and pulp	Ring plus a conjugated side chain extending delocalization	13–26
Flavan-3-ols	(+)-Catechin, (-)-epicatechin, epicatechin gallate	Seed mainly	Catechol B ring; saturated C ring blocks cross-ring conjugation	21–31
Proanthocyanidins	Procyanidins B1–B4, condensed tannins	Seed and skin	Two or more flavan-3-ol units; donor count multiplies	42 and above
Flavonols	Quercetin, myricetin, kaempferol and glycosides	Skin	2,3 double bond conjugated with 4-oxo; catechol B ring	22–43
Anthocyanins	Malvidin, peonidin, delphinidin 3-O-glucosides	Skin of red cultivars	Flavylium cation; charge state depends on pH	33–50
Stilbenes	trans-Resveratrol, piceid, piceatannol	Skin and stem	Two rings bridged by a conjugated double bond	17–40

II. EXTRACTION AND IDENTIFICATION

The most reported solid-liquid extraction (SLE) pathway. Milled pomace is suspended in a solvent, usually aqueous ethanol or methanol, sometimes at elevated temperature. It is inexpensive, general but the yield strongly depends on the solvent composition, temperature, contact time and solid to solvent ratio, and phenolics are degraded by long contact at high temperature [8]. Ethanol is preferred to methanol for food-facing work, as it is safer, easier to recover and already present in wine [1] Ultrasound-assisted extraction (UAE) uses acoustic cavitation to collapse bubbles against the cell wall, opening up the matrix and allowing solvent penetration. It shortens extraction time and is often reported to provide the richest list of phenolics resolved individually in chromatographic profiling. The same cavitation produces local heat and radicals, so sensitive compounds can be lost and reproducibility is affected by probe geometry and power. Microwave-assisted extraction (MAE) uses dielectric absorption to heat the solvent within the matrix, promoting rapid mass transfer. Reported effective for tannins and anthocyanins. The risk is also the rapid heating: anthocyanins have degraded if

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

power and duration were not controlled and the solvent has to absorb microwaves.

Supercritical fluid extraction (SFE) is based on the use of carbon dioxide above its critical point, usually with ethanol as a co-solvent. It is solvent-free and reported to be effective for the recovery of total polyphenols and flavonoids, but neat carbon dioxide is a poor solvent for polar polyphenols, the capital cost is high and yields shift sharply with pressure and temperature. Moutinho et al. [2]. optimized the concentration of ethanol for a single pomace against six endpoints and obtained six different optima: 69.6% for total phenolics, 55.1% for flavonoids, 90% for ortho-diphenols, 72.1% for ABTS, 53.0% for DPPH and 65.8% for FRAP. The best condition depends on what is being measured so no method or setting is universally better and Table 1 should be read as a set of tendencies. Drying is followed by extraction, and drying changes what survives. Oven drying is cheap, but exposes anthocyanins and stilbenes to heat for hours. For this reason, Tournour et al. [1] dried at 55 °C without forced air and ground under light protection. Freeze drying is more expensive, but yields a more friable material and better preserves labile compounds. Particle size matters too, since finer grinding increases surface area but risks thermal stress during milling. For modelling, this means a compound missing from a published profile may have been destroyed during drying rather than absent from the material, so absence from one study is weak evidence of absence in the cultivar. Cultivar surveys show how much the profile moves: five Sicilian reds [9], four Salento monovarietal [10] and four Romanian white and red varieties [11] each report a different dominant profile from the same tissue. High-performance liquid chromatography separates the extract in time; electrospray ionization puts the eluting compounds into the gas phase as ions; and tandem mass spectrometry fragments each ion and reads the fragments. The retention time, the ultraviolet spectrum from a diode-array detector and the fragmentation pattern together identify the compound, and a calibration curve against an authentic standard quantifies it.

Reported compound counts vary widely with the depth of the method and the sample. Moutinho et al. [2] identified 14 compounds across phenolic acids, flavonols and anthocyanins in two Portuguese red cultivars. Zhu et al. [3] identified 32 compounds in grape and kiwi juice pomace. Di Lorenzo et al. [12] reported 25 anthocyanins, 24 flavan-3-ols, 10 flavonols and 3 organic acids and stilbenes across a set of table and wine grapes. Counts in the region of 30 for pomace and higher for whole grapes are therefore normal, and the number is set as much by the analytical method as by the material. Five compounds are used throughout as worked examples: gallic acid, trans-resveratrol, (+)-catechin, quercetin and procyanidin B2. They were selected for structural spread. They include a hydroxybenzoic acid, a stilbene, a flavan-3-ol monomer, a flavonol and a flavan-3-ol dimer, with sizes ranging from 12 to 42 heavy atoms and phenolic hydroxyl counts from 3 to 8. That generalization makes them a test of whether structure can account for differences in activity, not a test of whether a model can recognize a single familiar scaffold. Abundance and intrinsic activity are distinct properties. Tournour et al. [1] recorded trans-resveratrol and quercetin below the detection limit but found gallic acid, caffeic acid, syringic acid, (+)-catechin and (-)-epicatechin by HPLC in Portuguese red pomace. A compound can be a good scavenger and a bad recovery target in a given cultivar and both facts should be reported.

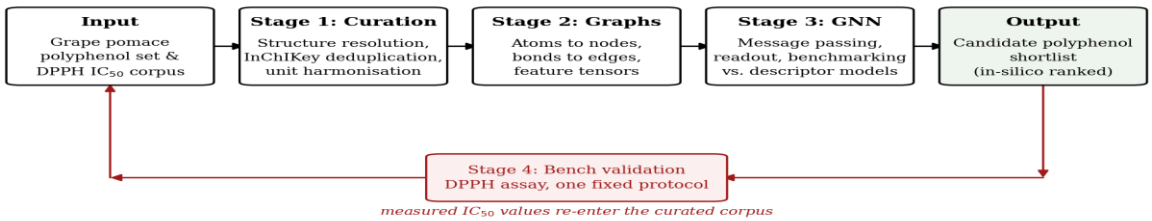

Fig. 1 Proposed workflow, from published pomace profiles through curation, graph construction, training and comparison to prediction and bench confirmation. Bench confirmation is a step inside the loop, feeding back into the data.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

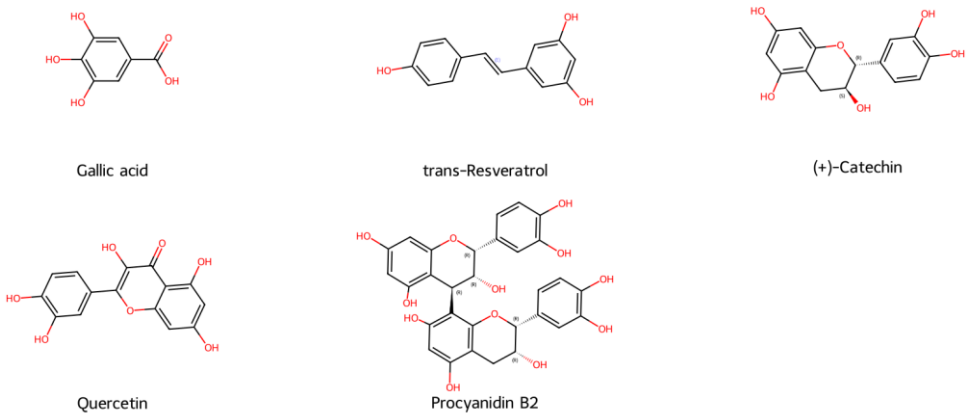

Fig. 2 The structures of the five target polyphenols, drawn from identifier-checked SMILES with stereocenters indicated.

Table 3
Properties of 5 target compounds. Here descriptors were computed from structures who's generated InChIKey matched the public reference value.

Property	Gallic acid	trans-Resveratrol	(+)-Catechin	Quercetin	Procyanidin B2
Class	Hydroxybenzoic acid	Stilbene	Flavan-3-ol	Flavonol	Proanthocyanidin
Molecular formula	C7H6O5	C14H12O3	C15H14O6	C15H10O7	C30H26O12
Molecular weight (g mol⁻¹)	170.12	228.25	290.27	302.24	578.53
CAS number	149-91-7	501-36-0	154-23-4	117-39-5	29106-49-8
Canonical SMILES	O=C(O)c1cc(O)c(O)c(O)c1	Oc1ccc(/C=C/c2cc(O)cc(O)c2)cc1	Oc1cc(O)c2c(c1)O[C@H](c1ccc(O)c(O)c1)[C@@H](O)C2	O=c1c(O)c(-c2ccc(O)c(O)c2)oc2cc(O)cc(O)c12	Oc1cc(O)c2c(c1)O[C@H](c1ccc(O)c(O)c1)[C@H](O)[C@H]2c1c(O)cc(O)c2c1O[C@H](c1ccc(O)c(O)c1)[C@H](O)C2
InChIKey (generated)	LNTHITQWFMADLM-UHFFFAOYSA-N	LUKBXSAWLPMMSZ-OWOJBTEDSA-N	PFTAWBLQPZVEMU-DZGCQCFKSA-N	REFJWTPEDVJJIY-UHFFFAOYSA-N	XFZJEEAOWLFHDH-NFJBMHMQSA-N
Matches reference key	Yes	Yes	Yes	Yes	Yes
Heavy atoms (graph nodes)	12	17	21	22	42
Bonds (graph edges)	12	18	23	24	47
Aromatic rings	1	2	2	3	4
Stereocentres	0	0	2	0	5
Phenolic hydroxyls	3	3	4	5	8
Catechol motifs	2	0	1	1	2

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Property	Gallic acid	trans-Resveratrol	(+)-Catechin	Quercetin	Procyanidin B2
Pyrogallol motifs	1	0	0	0	0
TPSA (Å²)	98	60.7	110.4	131.4	220.8
Calculated log P	0.5	2.97	1.55	1.99	3

Gallic acid returns two catechol matches because the two overlapping adjacent hydroxyl pairs inside its pyrogallol group are both counted.

III. ANTIOXIDANT ACTIVITY AND STRUCTURE

A free radical is a species with an unpaired electron. Reactive oxygen species are formed during normal metabolism and are increased by pollution, smoking and diet; in excess they attack lipids, proteins and DNA. A phenolic antioxidant quenches a radical by giving it either a hydrogen atom or an electron, and the antioxidant becomes a radical itself in the process. What makes phenolics useful is that the radical they become is comparatively unreactive. The unpaired electron does not stay on the oxygen; it spreads over the aromatic ring and any conjugated system attached to it. That spreading is resonance stabilization, and it is the reason a phenoxyl radical does not simply pass the damage along [13]. DPPH is a stable nitrogen-centered radical that is deep violet in solution and absorbs at 517 nm [14]. When an antioxidant reduces it, the colour fades, and the fall in absorbance is the signal [15, 16]. Fig. 3 sets out the reaction and how the result is read. The reaction can proceed by hydrogen atom transfer, in which the antioxidant hands over a hydrogen atom directly, or by electron transfer followed by proton loss. Which route dominates depends on the compound, the solvent and the pH, which is why DPPH is classed as a mixed-mechanism assay [17]. The solvent choice changes the measured value for the same compound and the magnitude of that shift is different among assays [18].

The results are usually given as IC50, the concentration needed to remove half the radical present. The lower the IC50, the better the scavenger. Standardization is important as the assay is stoichiometric and not catalytic: the IC50 is dependent on how much DPPH was in the well, the solvent, and when the plate was read as many phenolics react slowly and have not plateaued at 15 min [19]. Two IC50 values measured under different protocols are not the same quantity. The size of that effect on real data is quantified in Section 7.

ABTS uses a different stable radical cation and also runs by mixed mechanisms. FRAP measures reduction of an iron(III) complex under acid conditions and CUPRAC the reduction of a copper(II) complex, both single-electron-transfer assays that are pH dependent. ORAC measures how long an antioxidant protects a fluorescent probe from peroxyl radicals and is a hydrogen-transfer assay. Metal-chelating assays measure something different again, the ability to bind iron(II) and remove it from radical-generating chemistry (Table 4). These do not agree with one another. On the same pomace extracts, Tournour et al. [1] found total phenolic content correlated with DPPH at $R = 0.944$, with ORAC at $R \leq 0.632$, and with chelating ability at $R \leq 0.263$. Di Lorenzo et al. [12] found the same split between DPPH and ORAC and gave the chemical reason: compounds that react quickly with peroxyl radicals may react slowly with DPPH, or not at all, because the bulky DPPH radical cannot reach the donor site.

The units are also different quantities. Trolox equivalents per gram, percent inhibition and IC_{50} are not interconvertible and no factor relates them across compounds. Pooling assays and rescaling produces a number that no single experiment could generate, so a model that fits it well has fitted an artefact. If more data are needed, the better route is multitasking learning with one output head per assay over a shared representation, which keeps the labels separate [20]. Four structural features recur in the structure-activity literature [13, 21]. Phenolic hydroxyls are the donor sites, and an aromatic hydroxyl gives up its hydrogen far more readily than an alcohol on a saturated carbon, so the count that matters is of phenolic hydroxyls rather than of all hydroxyls. Aromatic rings accept the unpaired electron once the hydrogen has gone, and conjugation between rings widens the region over which it can spread. An ortho-dihydroxy, or catechol, arrangement is the most consistent of the four markers: the second hydroxyl stabilizes the radical formed on the first, both by hydrogen bonding and by extending delocalization (Fig. 3c), and three

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

adjacent hydroxyls in a pyrogallol or galloyl group do the same more strongly. Size and substitution modulate the rest, since bulky substituents can block access to the donor site and glycosylation of a hydroxyl removes it as a donor. For flavonoids specifically, three features are conventionally described as jointly required for high activity: a catechol B ring, a 2,3 double bond conjugated with the 4-oxo group, and hydroxyls at both position 3 and position 5 [21].

Table 3 and Fig. 8 give the computed descriptors. The five present very different environments to a radical. Gallic acid is a single ring carrying a pyrogallol group and a carboxyl, the smallest graph in the set. trans-Resveratrol has two rings bridged by a conjugated double bond and three hydroxyls, none of them adjacent. (+)-Catechin has a catechol B ring but a saturated C ring, so the B-ring radical cannot delocalize into the A ring. Quercetin has the same catechol B ring but is fully conjugated across all three rings and carries the 4-oxo group. Procyanidin B2 is two epicatechin units joined by an interflavan bond, doubling the donor count. One comparison shows why a structural model is worth trying. Both gallic acid and trans-resveratrol contain three phenolic hydroxyls. In gallic acid all three are on the same ring and adjacent to each other. In resveratrol they are spread over two rings with no adjacent pair. One that sees counts only treats them as equivalent, while one that sees the graph can distinguish between them. In this paper, we test whether such difference can be learned from data.

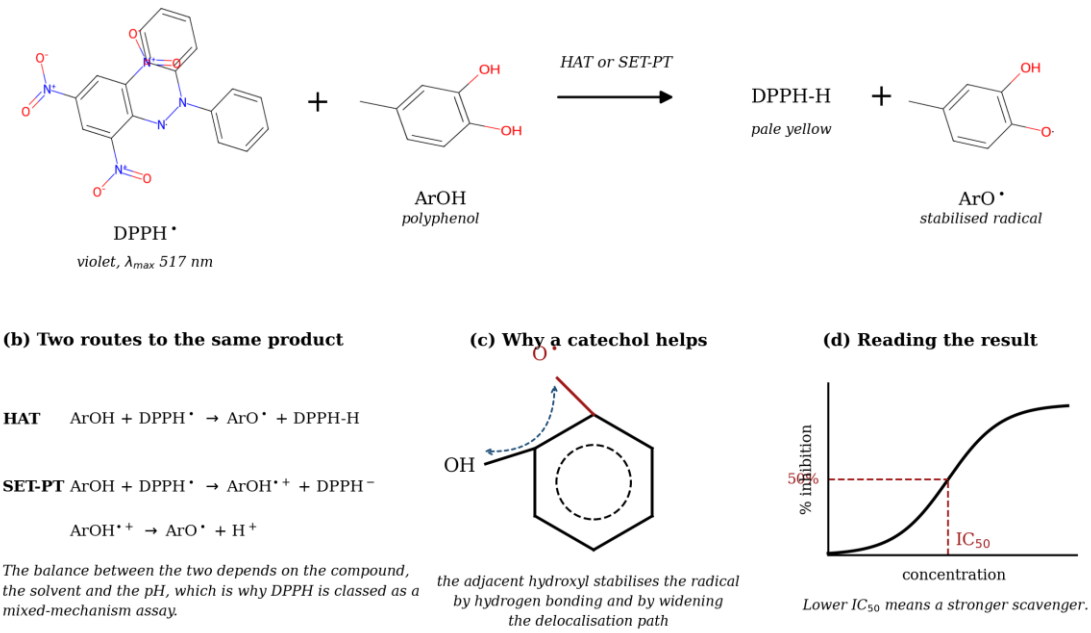

Fig. 3. The DPPH assay. (a) the reaction; loss of the violet colour at 517 nm is the signal. (b) the two routes to the same product. (c) why an adjacent hydroxyl stabilises the resulting radical. (d) how IC_{50} is read from a dose-response curve.

Table 4
Antioxidant assay families and their admissibility to a single-endpoint model.

Family	Reaction measured	Examples	Typical unit	Admitted
Hydrogen atom transfer	Donation of a hydrogen atom to a peroxyl radical	ORAC, TRAP, TOSC	Trolox equivalents; area under curve	No
Single electron transfer	Reduction of a metal ion or reagent; pH dependent	FRAP, CUPRAC, Folin-Ciocalteu	Trolox or gallic acid equivalents	No
Mixed	Both routes, in proportions set by compound, solvent and pH	DPPH, ABTS, DMPD	IC_{50}, SC_{50}, percent inhibition	DPPH IC_{50} only

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Family	Reaction measured	Examples	Typical unit	Admitted
Metal chelation	Binding of iron(II) or copper(II)	ICA, ferrozine assays	Percent inhibition per mg	No
Cellular	Effect inside a biological system	Cellular antioxidant activity	Various	No

On the same pomace extracts, Tournour et al. [1] found total phenolic content correlated with DPPH at R = 0.944, with ORAC at R ≤ 0.632 and with chelating ability at R ≤ 0.263.

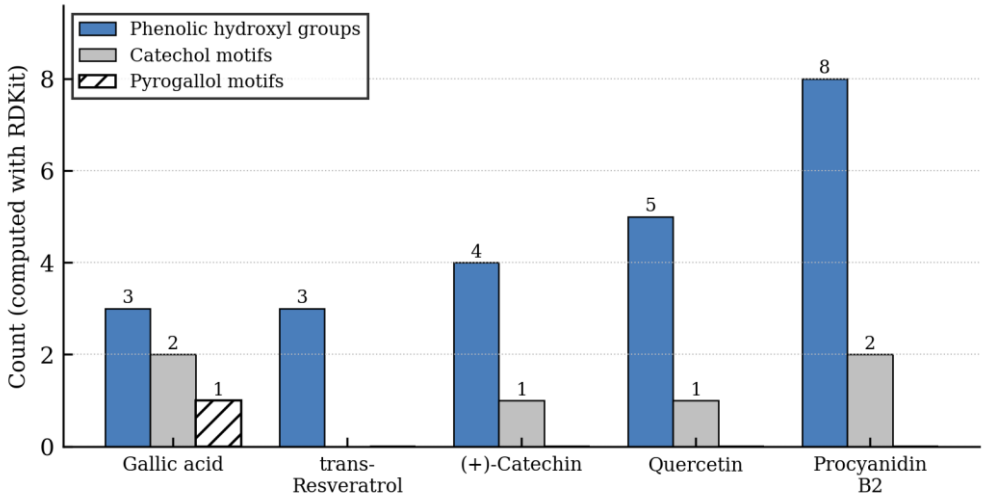

Fig. 4. Donor and motif counts for the five target compounds, computed from verified structures. These are counts, and they do not rank activity.

IV. FROM DESCRIPTORS TO GRAPHS

A quantitative structure-activity relationship model runs structure to descriptors to model to activity. A molecule is converted into a fixed vector of computed numbers, such as molecular weight, polar surface area, hydrogen-bond donor count or topological indices, and a regression is fitted from that vector to the measured property. Descriptor sets such as Mordred provide well over a thousand such numbers [22], usually cut down by feature selection before fitting. Tree ensembles are the dominant published work on this endpoint. Ghironi et al. [23] compared eleven algorithms on curated DPPH data, finding Extra Trees to be the best, Gradient Boosting and XGBoost close behind and linear models much worse. Kim et al. [24] found a similar picture with gradient boosting. Goya Jorge et al. [25] used DRAGON descriptors with a multilayer perceptron. Other methods have combined quantum chemical chemistry and classification to flag reactive hydrogens [26], and machine learning has been used to predict whether antioxidant combinations are synergistic or antagonistic [27].

Descriptors encode chemistry. In a different way. The limitation is that a person decides in advance which numbers to compute. If the feature that governs the chemistry is not in the set, the model cannot recover it, and connectivity is represented only indirectly, through indices that summarise the whole molecule rather than through the bonds themselves. Feature selection adds a second human decision on top of the first. One finding from the published DPPH models is still useful here. Ghironi et al. [23] reported that the most important descriptors in their best model related to electronic charge distribution and molecular geometry, which are the properties a graph encodes directly.

V. GRAPH NEURAL NETWORKS

A molecule is already a network: atoms joined to atoms by bonds. A graph neural network treats it as one. Atoms become nodes, bonds become edges, and the molecule becomes a graph carrying a single target value [28, 29]. Each node carries a feature vector describing one atom: element, formal charge, aromaticity, hybridisation, hydrogen count, degree,

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

ring membership and, optionally, an explicit donor-site flag. Each edge carries a vector describing one bond: bond order, conjugation, ring membership and stereochemistry. Table 6 gives the full specification with the reason for each entry. Hydrogens are not separate nodes; each heavy atom carries its hydrogen count, which keeps gallic acid at 12 nodes rather than 18. Fig. 4 shows the conversion and Fig. 5 the five target compounds as graphs. At the start each atom knows only about itself. In one round of message passing, every atom collects a message from each bonded neighbour, combines those messages with a function that does not depend on their order, and updates its own state. After one round an atom knows its neighbors; after three it knows the chemistry within three bonds. Order independence is required because a molecule has no natural atom numbering, and the same molecule written two ways must give the same answer. Depth is limited by over-smoothing. With too many rounds every atom has heard from every other, all node states converge, and the graph loses its internal structure. Two to four rounds is the usual range, and three was selected by grid search in published antioxidant graph work [30]. After message passing there is one vector per atom, but the prediction is one number per molecule. The readout collapses the atom vectors into a single molecular vector, the graph embedding, using a permutation-invariant operation such as mean, sum or maximum pooling. The embedding is the learned equivalent of a molecular fingerprint, except that it was learned for this task rather than defined in advance. The embedding passes through a small feed-forward network ending in one output unit with no activation, because pIC_{50} is unbounded. The architecture is shown in Fig. 6 and the hyperparameters in Table 7.

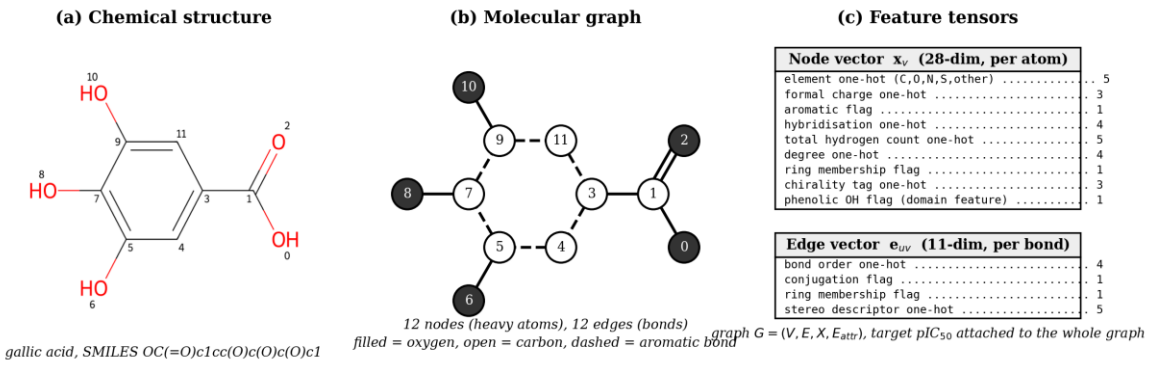

Fig. 5. Turning gallic acid into a molecular graph. (a) structure with atom indices; (b) 12 nodes and 12 edges, oxygen filled, aromatic bonds dashed; (c) the feature vectors handed to the network.

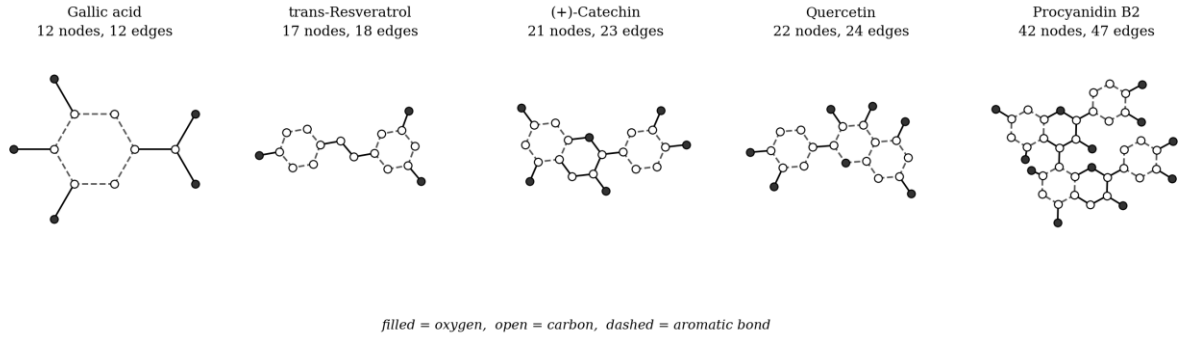

Fig. 6. The five target compounds as molecular graphs, from 12 nodes for gallic acid to 42 for procyanidin B2. One shared set of weights must handle the whole range.

Table 5
Node and edge features for the molecular graph, with the chemical reason for each.

Level	Feature	Encoding	Dim.	Why it matters here
Node	Element	One-hot: C, O, N, S, other	5	Oxygen is the donor site, the network must find them

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Level	Feature	Encoding	Dim.	Why it matters here
Node	Formal charge	One-hot: −1, 0, +1	3	Charge on the anthocyanins flavylium ion
Node	Aromaticity	Binary	1	A hydroxyl on an aromatic ring is a phenol, not an alcohol
Node	Hybridisation	One-hot: sp, sp2, sp3, other	4	Separates conjugated C ring of quercetin and the saturated one of catechin
Node	Hydrogen count	One-hot: 0–4	5	An oxygen without hydrogen cannot give one.
Node	Degree	One-hot: 1–4	4	Distinguishes a terminal hydroxyl from a ring ether oxygen
Node	Ring membership	Binary	1	Chain atoms are different from ring atoms
Node	Phenolic OH flag	Binary, optional	1	Not assumed but tested by ablation; explicit cue of donor site
Edge	Bond order	One-hot: single, double, triple, aromatic	4	Distinguishes the 2,3 double bond of quercetin from the single bond in catechin
Edge	Conjugation	Binary	1	conjugation is the way of the delocalizing radical
Edge	Ring membership	Binary	1	Ring bonds restrict geometry
Edge	Stereochemistry	One-hot: none, E, Z, cis, trans	5	The identity of resveratrol includes its trans form.

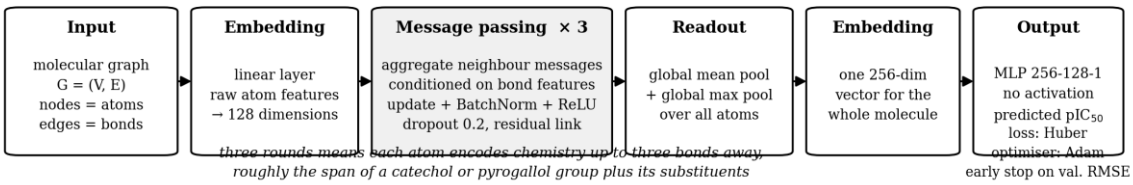

Fig. 7. Network architecture as used in Section 7. The output carries no activation function, because pIC_{50} is unbounded.

VI. MATERIALS AND METHODS

Fig. 1 shows the proposed workflow. It runs from published pomace profiles through compound selection, dataset construction, curation, graph building, training, validation, model comparison and application, ending in bench confirmation. Bench measurement is a step inside the loop, and its results feed back into the dataset. The reference dataset is built from the AODB antioxidant database [31] supplemented by primary literature. Records are included only if the assay is DPPH, the quantity is a quantitative IC_{50}, the molecular structure is available and unambiguous, and the units are convertible to a molar basis. Records are excluded if the structure is missing or cannot be parsed, the compound identity is ambiguous, the reaction time is unstated and cannot be recovered from the source paper, the unit is not convertible, or the value is a bound rather than a measurement. Table 5 shows the rule applying to each reported unit. Reaction time is set to a single window, as it is the biggest controllable source of variation between studies. We begin with 3790 DPPH entries according to Ghironi et al. [23], resolve

146 missing timing entries by reading the source publications, and retain 2832 compounds before deduplication (30 min window). Structures are parsed, salts and counterions removed, structures reduced to the largest organic fragment and canonized, and converted to a standard InChIKey [32, 33]. The generated key is compared against a public reference key before the structure is used.

The check catches real errors. A structure written here for (+)-catechin generated PFTAWBLQPZVEMU-UKRRQHHQSA-N, the identifier of (-)-epicatechin. The two differ only at carbon 3, and their formula, molecular weight and every count-based descriptor are identical, so a workflow keyed on names would have swapped them without warning. The corrected structure gives PFTAWBLQPZVEMU-DZGCQCFKSA-N, matching the reference. Because the identifier has separate blocks for connectivity and stereochemistry, a first-block match with a second-block mismatch points straight at the error. Duplicates are grouped by identifier, never by name, since the same molecule appears as catechin, (+)-catechin, D-catechin and cianidanol. For each group the coefficient of variation is computed; groups above 0.10 are discarded whole and logged, and groups at or below it are collapsed to the mean. Measurements must agree before they are merged. Every removed record goes to an exclusion log published as supplementary material.

Values are converted to mol L^{-1} by using the mass of the verified structure. Then they are transformed using the formula:

$$pIC50 = -\log_{10}(IC50\ [, \text{mol } L^{-1}]) \quad (1)$$

Raw IC50 values cover orders of magnitude and are highly skewed. The transformation helps make the distribution more symmetric. This is the reason given in the published curation of this endpoint [23]. It also flips the direction so that higher means stronger. Both the original value with its unit and the transformed value are kept so every conversion can be audited. Each standardized structure becomes a graph G = (V, E) with V the heavy atoms and E the bonds between them. Node and edge feature vectors follow Table 6, giving 28 and 11 dimensions respectively. Edges are stored in both directions because the graph is undirected. The target is attached to the whole graph, making this graph-level regression.

The architecture is a message-passing neural network with edge conditioning, so that the message a neighbour sends depends on the bond connecting it. Data are split into training, validation and independent test partitions. The validation set selects hyperparameters and the stopping point; the test set is used once, at the end, and informs no decision. Hyperparameters are tuned by nested cross-validation with the search confined to an inner loop, over the ranges in Table 7. Training repeats across at least five random seeds, and results are reported as mean and standard deviation rather than a single run.

Four baselines run on identical splits with identical seeds: Random Forest, Extra Trees, XGBoost and a feed-forward network on descriptors, plus a plain graph convolutional network as the simplest graph model [34]. Two further variants are alternatives to the message-passing network rather than baselines: graph attention networks weight neighbours by learned importance during message passing [35], and graph isomorphism networks use sum aggregation followed by a multilayer perceptron to maximize discriminative power [36]. Inputs to the first four are Morgan fingerprints at radius 2 with 2048 bits together with a descriptor block. Extra Trees is included because it holds the best published result on this endpoint. The descriptor network is the baseline most often left out, and it is the one that makes the comparison readable, because without it a graph win could mean either that the graph helped or that a neural network helped. Three metrics are reported together, because each hide something. R^2 is the share of variance the model accounts for, but it depends on the spread of the test set and flatters a weak model on a wide one [37]. MAE is the average error size and treats all errors alike, so a few very bad predictions can be masked. RMSE penalizes large errors, and the gap between RMSE and MAE shows how uneven the errors are. All are in log units, so an MAE of 0.31 means the typical prediction is out by about a factor of two in concentration. Predicted-against-measured scatter plots, residual plots and training and validation loss curves should be reported alongside.

Splitting is by Bemis-Murcko scaffold [38] rather than at random. Chemical datasets contain families of close analogues, so a random split nearly always leaves a relative of every test molecule in training and the score reflects interpolation rather than generalization. A scaffold split assigns whole ring-system families, so no test scaffold appears in

training. It gives lower numbers, and those lower numbers are the more honest estimate. Random-split results are reported alongside because the published benchmarks used random splitting. The remaining safeguards are: deduplication before splitting, not after; label scrambling, where targets are shuffled and the same pipeline retrained to confirm the score collapses towards zero; an applicability domain by bounding box and distance to centroid, following the third OECD principle [39]; the maximum similarity of each test molecule to the training set, reported as a distribution; an uncertainty interval on every prediction; and an external test set of compounds published after the cut-off or measured in house.

The frozen model is tested against the five target compounds. For each of them, the study reports structure, class, molecular formula, canonical SMILES, molecular weight, reported occurrence in pomace, experimental antioxidant value available in literature and model output. Any target compound already present in the training set should be identified and reported as such; gallic acid and quercetin are standard reference compounds in DPPH work and are likely to be there, and for those the output to report is the measured value from the dataset rather than a prediction. Table 13 gives the template.

Structure handling, descriptors and graph construction use RDKit [40]. Baselines use scikit-learn and XGBoost, graph models PyTorch Geometric. Versions are recorded in the methods, because descriptor values, and calculated partition coefficients in particular, differ between toolkits and between versions.

Table 6
Architecture and hyperparameters as used in the run of Section 7. Values were fixed in advance and not tuned, because tuning against a 16-compound cross-validation score would be leakage.

Parameter	Value used	Reason
Architecture	Message-passing network with edge conditioning	A bond type should change the message it carries [28]
Message-passing layers	3	Approximately the reach needed to sense a catechol or pyrogallol group; also selected by grid search in [30]
Hidden dimension	48	Scaled down from the usual 128 for a 16-compound corpus
Aggregation	Sum	Preserves neighbour counts, which matter when donor count matters
Readout	Mean and max pooling concatenated	Sum pooling makes the vector grow with atom count, so the model could predict from size
Dropout; weight decay	0.15; 8×10^{-4}	Regularization for a very small corpus
Loss; optimizer; learning rate	Huber; Adam; 6×10^{-3}	Tolerant of residual outliers; standard practice
Epochs	350, full-batch	Corpus fits in one batch
Validation	Leave-one-out and scaffold-grouped 4-fold	n = 16 leaves no room for a held-out test set
Ensemble	6 seeds for target prediction	Provides the uncertainty interval in Table 13

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

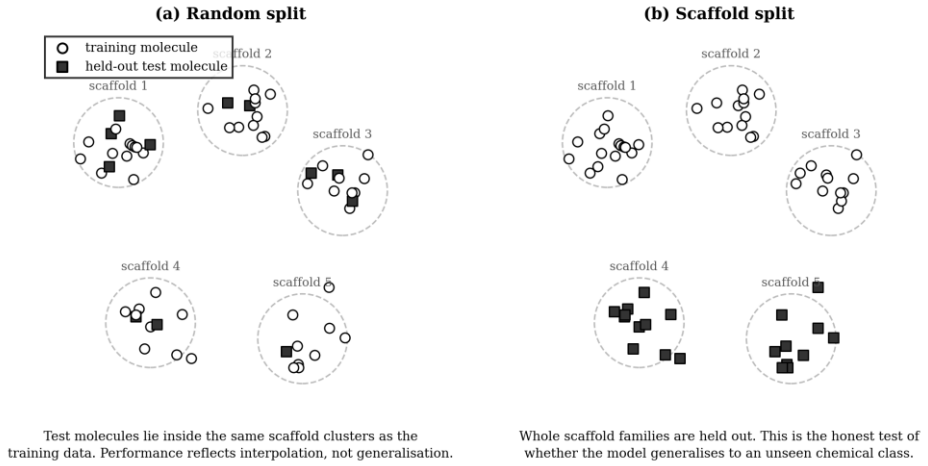

Fig. 8. Random splitting flatters a molecular model and scaffold splitting does not. The gap between the two is itself worth reporting.

VII. RESULTS

This section reports what the pipeline produced. It describes the target compounds and the measured data available for them, the corpus collected for training, the model that was adapted to the corpus, and the predictions provided by the model for the five grape pomace polyphenols. All below are either published measurements, calculations from a verified structure or an output of the model trained here. Identity check and calculated descriptors for the 5 target compounds are reported in Table 3. Donor and motif counts are plotted in Fig. 8. All five generated InChIKeys matched the respective public reference values. Gallic acid gives two catechol matches and one pyrogallol match because both overlapping neighboring hydroxyl pairs within its pyrogallol group are counted. trans-Resveratrol has no motif, but it has three phenolic hydroxyls. (+)-Catechin and quercetin both give back one catechol on the B ring but are different in that quercetin is conjugated on all three rings. Procyanidin B2 has the highest donor count of eight donors and two catechols, one per epicatechin unit. There are two conclusions. Donor count and donor adjacency are independent structural properties, so a count-based representation is unable to differentiate gallic acid from resveratrol. And procyanidin B2 has the most donors and the highest molecular weight, so a model whose predictions correlate with molecular weight would rank it highly for reasons that have nothing to do with chemistry; the correlation of prediction and molecular weight is therefore reported as a diagnostic. The donor column is not a potency ranking. Bond dissociation enthalpy at the weakest hydroxyl, solvent effects, steric access and reaction kinetics all count and none appear in the table.

Fig. Figure 4 shows the conversion for gallic acid. 12 heavy atoms give 12 nodes and 12 bonds give 12 edges, but these are held as 24 directed edges in the implementation. Atoms 6, 8 and 10 are the phenolic oxygens. Fig. 5 shows all five compounds as graphs, from 12 nodes for gallic acid to 42 for procyanidin B2, which is the size range one shared set of weights must handle. The features in Table 6 were chosen with this chemistry in mind. The aromaticity flag exists because a hydroxyl on an aromatic ring is a phenol and one on a saturated carbon is an alcohol. The hydrogen count exists because an oxygen with no hydrogen cannot donate one. The bond conjugation flag exists because conjugation is the path the radical delocalizes along.

Before any model was built, published pure-compound DPPH IC_{50} values were recovered for the target compounds and converted to a molar basis by Eq. (1). Table 10 gives the records and Fig. 10 the resulting picture. Two studies supply most of the set: Iacopini et al. [41] measured catechin, epicatechin, quercetin, rutin and resveratrol as pure standards, and Lv et al. [42] measured six procyanidin standards including procyanidin B2 against a 25µg mL^{-1} DPPH solution read at 60 min. (-)-Epicatechin was measured in both studies, which makes it a bridge between them. Iacopini et al. [41] report 6.8 µM; Lv et al. [42] report 1.88 µg mL^{-1}, which for a molar mass of 290.27 g mol^{-1} is 6.48 µM. The two differ by 4.8 percent in concentration, or 0.021 pIC_{50} log units. Two laboratories, two matrices and two unit systems agree once the units are converted, which shows that

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

the conversion step of Section 6 matters. Another example is in the target set. Quercetin is mentioned as 5.59 µg mL-1 in a study on grape pomace [43] and 18.61 µM in another study. They seem like different answers until you convert them: 5.59 µg mL^{-1} divided by 302.24 g mol^{-1} is 18.50 µM, a 0.6 percent difference. The mass unit concealed the agreement.

Table 7
The literature values of pure compound DPPH IC50 for the five target compounds were extracted and converted to a molar basis by Eq. (1). Each entry is a measurement published by the authors cited.

Compound	Reported	Unit	→ µM	pIC_{50}	Protocol note	Source
trans-Resveratrol	81.92 ± 9.17	µM	81.92	4.087	Pure standard	Sy et al. [44]
Gallic acid	30.0 ± 2.9	µM	30.00	4.523	400 µM DPPH, assay control	US Patent 8,846,757 [45]
(-)-Epicatechin	6.8	µM	6.80	5.167	Pure standard	Iacopini et al. [41]
(-)-Epicatechin	1.88 ± 0.01	µg mL^{-1}	6.48	5.189	25 µg mL^{-1} DPPH, 60 min	Lv et al. [42]
(+)-Catechin	6.7	µM	6.70	5.174	Pure standard	Iacopini et al. [41]
Quercetin	5.5	µM	5.50	5.260	Pure standard	Iacopini et al. [41]
Quercetin	5.59 ± 0.13	µg mL^{-1}	18.50	4.733	Assay standard; 100 µM DPPH, 30 min	Luchian et al. [43]
Quercetin	18.61 ± 3.55	µM	18.61	4.730	Pure standard	Not traced; citation to be supplied
Procyanidin B2	1.88–2.70	µg mL^{-1}	3.25–4.67	5.331–5.488	25 µg mL^{-1} DPPH, 60 min	Lv et al. [42], bounded
Ascorbic acid	2.38 ± 0.03	µg mL^{-1}	13.51	4.869	Reference compound	Lv et al. [42]

(-)-Epicatechin was quantified in both primary studies and is the cross-anchor: 6.80 and 6.48 µM, a difference of 0.021 pIC50 log units. Procyanidin B2 is bounded. The source reports a range across six standards. These records are not used for training they give the scale and give an external check on the model.

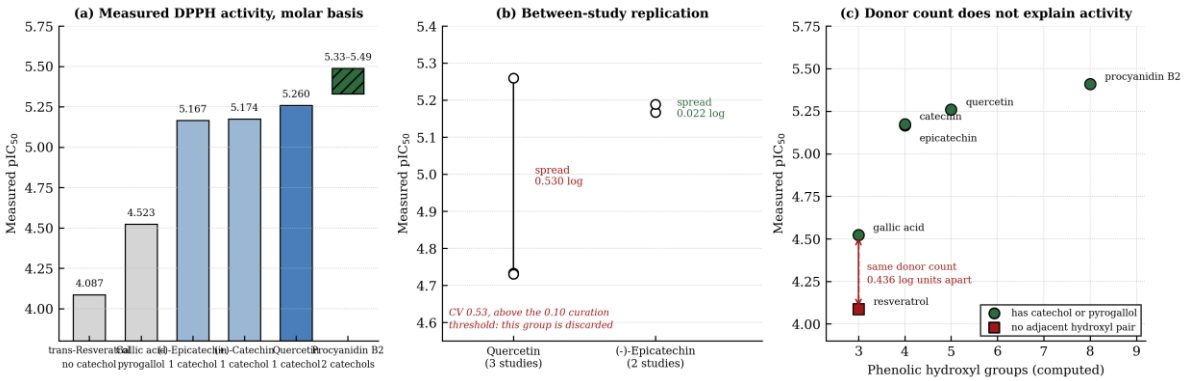

Fig. 10. Measured DPPH activity of five target compounds Collated from the sources in Table 10 (a) pIC50 on a molar basis with the radical-stabilizing motif annotated. (b) Between-study replication: epicatechin replicates to 0.021 log units across two laboratories and quercetin to 0.529 across three. (c) Activity vs donor count. Gallic acid and trans-resveratrol have the same three hydroxyls and differ by 0.436 log units.

The best-documented pIC50 order is procyanidin B2 (5.33 to 5.49) > quercetin (5.260) > catechin (5.174) ≈ epicatechin (5.167) > gallic acid (4.523) > trans-resveratrol (4.087). Procyanidin B2 is indicated as a bounded interval since Lv et al. [42] give a range across their six standards, rather than a single value. The measurements lead to three structural readings. Gallic acid and trans-resveratrol have the same three phenolic hydroxyls, but differ by 0.436 log units, or a factor of 2.73. Gallic acid has a pyrogallol group and resveratrol has no adjacent hydroxyl pair anywhere. Quercetin and catechin both have one B-ring catechol, and quercetin adds full conjugation and a 4-oxo group for only a 0.086 log unit gain, so the catechol does most of the work. The dimer, procyanidin B2, has double the donors of its monomer, but only gains about 0.24 log units, a factor of 1.75 rather than 2. Per unit mass the dimer and monomer are almost indistinguishable. Replication varies widely by compound, setting a lower bound on what any fitting model to these numbers can resolve

Epicatechin (measured twice) varies 0.021 log units, coefficient of variation 0.034. Quercetin, measured three times, ranges from 0.529 log units with a coefficient of variation of 0.531. The target set spans a total of 1.32 log units and the signal-to-noise ratio relative to the worst replica is around 2.5. Quercetin is the most well-documented compound in the set and under the reconciliation rule of Section 6 the quercetin group would be discarded entirely. In principle a threshold of 0.10 is defensible but in practice severe, and any study adopting it should report what proportion of its corpus it drops. To get a large enough corpus to train on, many compounds under one protocol had to be measured from a source. Moazzen et al. [46] measured

DPPH IC_{50} for 25 phenolic compounds spanning hydroxybenzoic acids, hydroxycinnamic acids, flavonoids and two synthetic antioxidants, all in one laboratory: 0.1 mmol L^{-1} DPPH in 50 mmol L^{-1} Tris-HCl at pH 7.4, 30 min in the dark, read at 517 nm. That window of 30 min is precisely the restriction of Section 6. Six compounds were reported as not calculated or not detected as a numeric value for IC50 and were removed: 2,4-dihydroxybenzoic acid, 2,6-dihydroxybenzoic acid, 4-hydroxybenzoic acid, quinic acid, cinnamic acid and o-coumaric acid. And Nineteen

Then, a second filter that was not anticipated by the Section 6 protocol was implemented. Each dose-response regression's coefficient of determination is reported with the IC50 by Moazzen et al. [46] allowing for a direct assessment of the significance of the fitted value. Three substances m-coumaric acid (0.188), n-propyl gallate (0.405), and rutin (0.677)—failed at a threshold of 0.70. The case of m-coumaric acid demonstrates the importance of the filter. Its tabulated IC_{50} of 3.328 μg mL^{-1} is the lowest in the entire table and would have made it the most potent compound in the corpus, yet the authors state in their own text that its activity was negligible. The number is an extrapolation from a regression explaining 19 percent of the variance. A curation protocol filtering only on units and assay type accepts it and places a spurious outlier at the top of the activity range. After InChIKey deduplication, which removed nothing further, the corpus contains 16 compounds spanning 1.477 pIC_{50} log units, from vanillic acid at 2.679 to 3,5-dihydroxybenzoic acid at 4.156, with a mean of 3.432 and a standard deviation of 0.432. Six distinct Bemis-Murcko scaffolds are represented and graph sizes run from 11 to 32 heavy atoms. Table 11 gives the corpus and Fig. 11(a) shows it. Three compounds in the corpus have published DPPH values from other laboratories, which permits a direct test of whether protocols can be pooled. They cannot. Gallic acid is 9.8 times weaker in the Moazzen protocol than in the comparison source, catechin 24.2 times weaker and BHT 12.1 times weaker. The mean offset is 1.152 pIC_{50} log units with a standard deviation of 0.206, shown in Fig. 11(b).

Table 8
The curated training corpus. All 16 compounds were measured by Moazzen et al. [46] under one protocol: 0.1 mmol L^{-1} DPPH in 50 mmol L^{-1} Tris-HCl pH 7.4, 30 min, 517 nm.

Compound	Class	Formula	MW	IC_{50} (μg mL^{-1})	IC_{50} (μM)	pIC_{50}	Fit R^2	InChIKey
3,5-Dihydroxybenzoic acid	hydroxybenzoic acid	C7H6O4	154.12	10.755	69.8	4.156	0.973	UYEMGAFJOZZIFP-UHFFFAOYSA-N

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Compound	Class	Formula	MW	IC_{50} (µg mL^{-1})	IC_{50} (µM)	pIC_{50}	Fit R^2	InChIKey
Rosmarinic acid	hydroxycinnamic acid	C18H16O8	360.32	30.288	84.1	4.075	0.856	DOUMFZQKYFQNTF-WUTVXBCWSA-N
(+)-Catechin	flavonoid	C15H14O6	290.27	47.138	162.4	3.789	0.996	PFTAWBLQPZVEMU-DZGCQCFKSA-N
Luteolin	flavonoid	C15H10O6	286.24	52.444	183.2	3.737	0.952	IQPNAANSBPBGFQ-UHFFFAOYSA-N
2,3-Dihydroxybenzoic acid	hydroxybenzoic acid	C7H6O4	154.12	29.559	191.8	3.717	0.866	GLDQAMYCGOIJDV-UHFFFAOYSA-N
Gentisic acid	hydroxybenzoic acid	C7H6O4	154.12	30.901	200.5	3.698	0.867	WXTMDXOMEHJXQO-UHFFFAOYSA-N
Caffeic acid	hydroxycinnamic acid	C9H8O4	180.16	49.382	274.1	3.562	0.979	QAIPRVGONGVQAS-DUXPYHPUSA-N
Gallic acid	hydroxybenzoic acid	C7H6O5	170.12	49.913	293.4	3.533	0.871	LNTHITQWFMADLM-UHFFFAOYSA-N
3,4-Dihydroxybenzoic acid	hydroxybenzoic acid	C7H6O4	154.12	47.524	308.4	3.511	0.973	YQUVCSBJEUQKSH-UHFFFAOYSA-N
Isoorientin	flavonoid	C21H20O11	448.38	191.933	428.1	3.368	0.894	ODBRNZZJSYPIDI-IAAKTDFRSA-N
Taxifolin	flavonoid	C15H12O7	304.25	231.440	760.7	3.119	0.877	CXQWRCVTCMQVQX-GJZGRUSLSA-N
BHT	synthetic	C15H24O	220.36	191.998	871.3	3.060	0.998	NLZUEZXRPGMBCV-UHFFFAOYSA-N
Syringic acid	hydroxybenzoic acid	C9H10O5	198.17	202.150	1020.1	2.991	0.928	JMSVCTWVEWCHDZ-UHFFFAOYSA-N
Myricetin	flavonoid	C15H10O8	318.24	345.848	1086.8	2.964	0.747	IKMDFBPHZNJCSN-UHFFFAOYSA-N
Ferulic acid	hydroxycinnamic acid	C10H10O4	194.19	217.730	1121.2	2.950	0.957	KSEBMYQBYZTDHS-HWKANZROSA-N
Vanillic acid	hydroxybenzoic acid	C8H8O4	168.15	351.993	2093.3	2.679	0.992	WKOLLVMJNQIZCI-UHFFFAOYSA-N

Fit R^2 is the coefficient of determination of the dose-response regression as reported by the original authors; records below 0.70 were excluded during curation (m-coumaric acid 0.188, n-propyl gallate 0.405, rutin 0.677). All 25 structures were resolved here and every molecular weight checked against the expected value for the named compound.

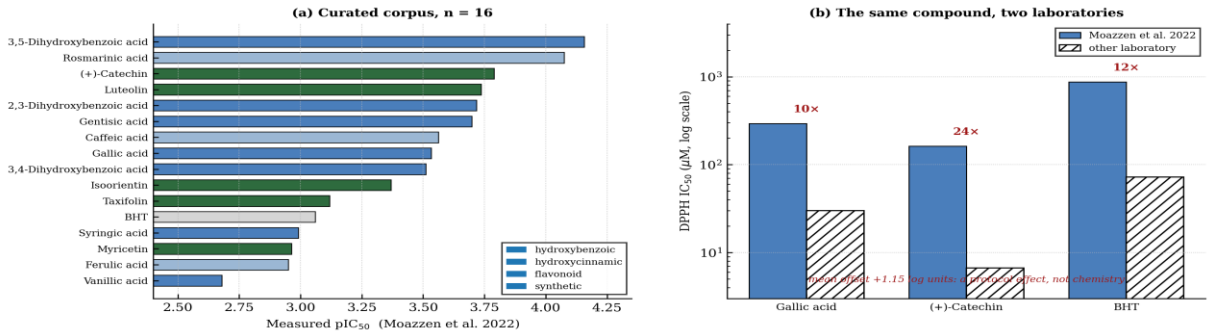

Fig. 9. The training corpus and why it was not enlarged. (a) The 16 curated compounds ranked by measured pIC_{50}, colored by structural class. (b) Three compounds with values from a second laboratory on a logarithmic axis; the offset is systematic at +1.152 log units, larger than the entire dynamic range of the corpus.

What matters is how consistent that offset is. A spread of 0.206 around a shift of 1.152 is a systematic protocol effect rather than compound-specific chemistry. Pooling would inject an artefact larger than the entire 1.477 log unit span of the corpus, and any model fitted to the merged set would spend capacity learning which laboratory produced each measurement. The corpus was therefore held to one protocol at the cost of a smaller n, which is the single-assay rule of Section 6 applied to real data. A message-passing network with 28-dimensional node features, three rounds of edge-conditioned message passing, concatenated mean and maximum readout and a two-layer head was trained against Random Forest, Extra Trees and cross-validated Ridge regression on Morgan fingerprints plus computed descriptors, together with a mean-predictor null model. Leave-one-out cross-validation was used because n = 16 leaves no room for a conventional split. Table 12 gives the results and Fig. 12 shows them. No model did meaningfully better than the null. The mean predictor achieves R^2 of −0.138. Random Forest reaches 0.013, Extra Trees −0.083, the graph network −0.744 and Ridge −0.721. A negative R^2 means the model is worse than predicting the corpus average for every compound, and the graph network is the second worst performer in the set. The reason is visible in the training fit. The network reaches R^2 of 0.696 on the data it was trained on and −0.744 when the compound is held out, a drop of 1.440. It is memorising the 16 molecules rather than learning chemistry from them. Fig. 12(b) shows the consequence: predictions cluster on the corpus mean and carry almost no gradient with the measured value. Scaffold-grouped cross-validation gave −0.796 against −0.744 for leave-one-out, a difference of only 0.053. That small gap says nothing good about the model; a model already performing at the level of the mean cannot be made much worse by a harder split.

The control specified in Section 6 tests whether that fit carries any information: shuffle the target values, retrain the identical pipeline, and check that performance collapses. Across five shuffles the scrambled model achieved leave-one-out R^2 values of −0.845, −0.851, −0.776, −1.826 and −0.525, a mean of −0.965. The best scrambled run, at −0.525, scored better than the model trained on real labels at −0.744, which Fig. 12(c) shows directly. At this sample size the network therefore recovers no structure-activity signal, and the run is a pilot that establishes the pipeline rather than a predictive model.

The frozen model was nonetheless applied to the five target compounds so the output can be examined for what it is. Gallic acid and catechin are in the corpus and are reported as leave-one-out predictions; quercetin, trans-resveratrol and procyanidin B2 are absent and are genuine out-of-sample predictions. Table 13 gives the values and Fig. 13 shows them with 95 percent ensemble intervals.

Table 9

Model performance on the 16-compound corpus. Leave-one-out cross-validation was used because the corpus is too small for a held-out test set. Errors are in pIC_{50} log units.

Model	Molecular input	R^2	RMSE	MAE
MPNN (molecular graph)	28-dim node + adjacency	-0.744	0.552	0.414
Random Forest	Morgan FP + descriptors	0.013	0.415	0.357
Extra Trees	Morgan FP + descriptors	-0.083	0.435	0.366
Ridge (cross-validated)	Morgan FP + descriptors	-0.721	0.549	0.458
Mean predictor (null model)	None	-0.138	0.446	0.386
MPNN, training-set fit	28-dim node + adjacency	0.696	0.230	0.193
MPNN, scaffold-grouped CV	28-dim node + adjacency	-0.796	0.560	0.501
MPNN, scrambled labels (best of 5)	28-dim node + adjacency	-0.525		

The first five rows are leave-one-out. The training-set fit exposes the gap of 1.440 R^2 units between memorisation and generalisation. The scrambled-label row is the best of five shuffles: at −0.525 it scores better than the model trained on true labels (−0.744), which is the criterion set in Section 6 for concluding that the model has recovered

no signal at this sample size. Published benchmarks on this endpoint reach 0.77 to 0.81 on corpora of 1911 and 3133 compounds (Table 8).

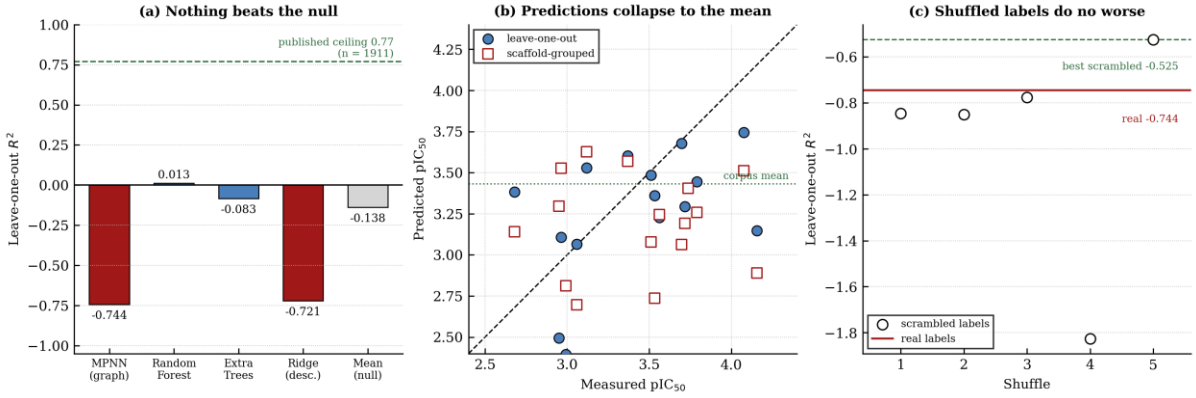

Fig. 10. Model results. (a) Leave-one-out R^2 for every model tested; the dashed line is the published ceiling at n = 1911. Nothing reaches the mean-predictor null at −0.138 except Random Forest at 0.013. (b) Predicted against measured pIC_{50} under both validation schemes; predictions cluster on the corpus mean. (c) The scrambling control: the best shuffle at −0.525 beats the real model at −0.744.

The predicted ranking is procyanidin B2 (3.522) > catechin (3.446) > trans-resveratrol (3.428) > gallic acid (3.360) > quercetin (3.200). The entire spread across five compounds spanning four structural classes is 0.322 log units, less than one standard deviation of the corpus and roughly a factor of two in concentration. All ten pairwise comparisons overlap at 95 percent, so the model separates none of the five. Where a measured value exists the residuals are −0.172 for gallic acid and −0.344 for catechin, both under-predictions.

The predicted ranking also contradicts the independent measurements reported earlier in this section, which place quercetin among the strongest scavengers in the set while the model places it last. Table 13 therefore says nothing about which grape pomace polyphenol is the better antioxidant. It is included because a reader shown only the ranking, without the scrambling control and the interval widths in Fig. 13, could take it for a result, and the reporting rules of Section 6 exist to prevent that. Table 8 reproduces the results of Ghironi et al. [23] for eleven regression algorithms on 1911 compounds filtered to DPPH IC_{50} under a 30 min protocol, and Fig. 7 shows the same data. Ensemble trees clearly beat linear models, and the ceiling is R^2 of 0.77 for Extra Trees and 0.78 for a consensus of the three best, with RMSE near 0.45 log units. Kim et al. [24] reported R^2 of 0.81 on 3133 records from the same database and Goya Jorge et al. [25] a training R^2 of 0.713 with an external Q^2 of 0.654 on 1373 compounds.

In particular, Cai et al. [30] tested five graph designs on three antioxidant benchmarks, and message passing prevailed in all three. Two of their conclusions hold true: the graph model improved as molecules got bigger, and the graph model by itself was not the best model because a hybrid that combined graph and sequence features outperformed both. In comparison to those numbers, the run presented here has a sample size that is two orders of magnitude smaller and, as a result, performs worse: 16 compounds yield nothing above the null, while 1911 yield 0.77. A comprehensive study on this endpoint would need to map the curve between those two places.

Table 10

The five grape pomace targets predicted DPPH activity. The training corpus contains gallic acid and catechin, which are reported as leave-one-out predictions; the remaining three are actual out-of-sample predictions. The standard deviation of a six-person ensemble is the measure of uncertainty.

Compound	Class	In corpus	Measured pIC_{50}	Predicted pIC_{50}	SD	Residual	Predicted IC_{50} (µM)
Procyanidin B2	proanthocyanidin	no	not measured	3.522	0.141	n/a	301

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Compound	Class	In corpus	Measured pIC_{50}	Predicted pIC_{50}	SD	Residual	Predicted IC_{50} (μM)
(+)-Catechin	flavan-3-ol	yes	3.789	3.446	0.144	-0.344	358
trans-Resveratrol	stilbene	no	not measured	3.428	0.199	n/a	373
Gallic acid	hydroxybenzoic acid	yes	3.533	3.360	0.103	-0.172	436
Quercetin	flavonol	no	not measured	3.200	0.102	n/a	630

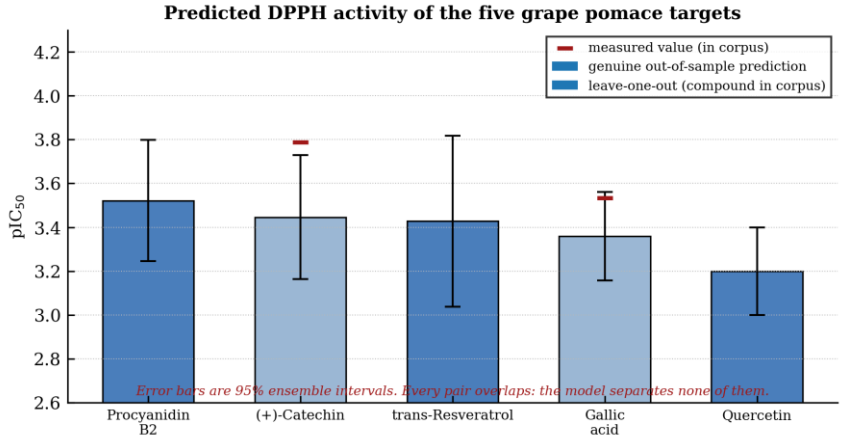

Fig. 11. Predicted DPPH activity of the five grape pomace targets with 95 percent ensemble intervals. Red dashes mark the measured value where the compound is in the corpus. All five intervals overlap one another.

VIII. DISCUSSION

The scrambling control is what makes the rest of the run interpretable. Without it, a leave-one-out R^2 of −0.744 could be put down to a hard endpoint, a difficult chemical space or an unlucky split, and could be improved by trying architectures until one produced a positive figure. With it the reading is settled: shuffled labels scored as well, so there is nothing at this sample size for an architecture search to find, and any positive figure obtained that way would be fitted noise. Running the control costs one experiment and a few minutes of compute, and it is what separates a pilot that can be built on from one that cannot. The measurements gathered along the way remain informative even though the model is not.

Within the hydroxybenzoic series, 3,5-dihydroxybenzoic acid is the most active compound in the corpus at pIC_{50} 4.156, ahead of gentisic acid, 3,4-dihydroxybenzoic acid and gallic acid. Moazzen et al. [46] interpret this as proof that hydroxylation position is more important than total count and observe that activity is decreased when a para hydroxyl is added to 3,5-dihydroxybenzoic acid to produce gallic acid. The two weakest chemicals in the corpus are vanillic and syringic acid, which have methoxy groups in place of hydroxyls. Myricetin is the weakest flavonoid and catechin is the strongest, which defies a straightforward donor-count prediction because myricetin has six hydroxyls overall and a pyrogallol B ring. Catechin also outperforms taxifolin, its close analogue differing only by a 4-oxo group. These are real structure-activity relationships and they are not simple ones: position matters more than count, an extra hydroxyl can reduce activity, and the flavonoid ordering runs opposite to the naive expectation. A representation that encodes connectivity should in principle capture all of this, and that it did not at n = 16 says more about the sample size than about the representation. What such a representation would show, were the corpus large enough to train on, is worth setting out in advance.

Attributions from a well-trained model should fall on the catechol and pyrogallol oxygens, on the conjugated systems attached to them, and not on the sugar residues of glycosides. Attention weights and gradient-based attributions can be inspected for this, and attention-based molecular architectures were

designed with that inspection in mind [47]. There is a firm limit on what such an analysis can claim. Attributions show which parts of the input the model relied on, and agreement with accepted chemistry raises confidence that the model has learned something real rather than a dataset artefact. They say nothing about the chemistry itself, and no model demonstrates a mechanism. The same caution applies to the output the framework is meant to produce.

That output is a priority list for testing. A defensible statement is that a model identifies a compound as a high-priority candidate for experimental testing; a claim that a model proves a compound is the strongest antioxidant is not defensible, and the present run does not support even the weaker statement. Scaffold analysis returns five distinct Bemis-Murcko scaffolds for the five target compounds: a benzene ring, a stilbene, a flavan, a flavone and a bis-flavan. The fact that none of them can be predicted by remembering a close sister is a positive characteristic, but it also indicates that the five are situated at significantly different distances from any training corpus. The bis-flavan of procyanidin B2 is complex and likely rare, whereas the benzene scaffold of gallic acid is the most prevalent in any phenolic combination. Since testing only, the compounds a model favors cannot identify false positives at the top or a systematic bias elsewhere, that distance should be determined prior to training and given per compound. Confirmation at the bench should sample over the projected range rather than just its top.

IX. CONCLUSIONS

Grape pomace polyphenols are individually under-measured, and extract-level results cannot be decomposed into compound-level ones because mixtures behave non-additively. Predicting DPPH activity from molecular structure is a reasonable way to decide what to assay first, and a graph representation suits the chemistry, since the groups that govern radical scavenging are neighbourhood features of the molecular graph. This paper set out that pipeline in full and then ran it. Three findings concern the data. Unit conversion matters: epicatechin measured in two laboratories in two unit systems agrees to 0.021 log units once converted, and quercetin reported as 5.59 μg mL^{-1} and as 18.61 μM is the same measurement twice. Donor count does not order the compounds, since gallic acid and trans-resveratrol carry the same three phenolic hydroxyls yet differ by 0.436 log units, which is the strongest available argument for a representation that encodes adjacency. And protocols cannot be pooled, since three compounds measured in two laboratories differ systematically by 1.152 log units, more than the dynamic range of the training corpus. On the modelling side, 25 published measurements made under one protocol reduced to 16 usable compounds once non-numeric entries were removed and dose-response fit quality was filtered on. On those 16 a message-passing network reached a leave-one-out R^2 of −0.744 against −0.138 for a mean predictor and −0.796 under scaffold grouping, with a training fit of 0.696 that identifies memorisation, and the label-scrambling control confirmed that no structure-activity signal was recovered at this sample size. The predicted ranking of the five grape pomace targets spans 0.322 log units, every pairwise comparison overlaps at 95 percent, and the ordering contradicts independent measurement, so it is reported only to show what the same run would look like without its controls. Two protocol amendments follow from the work: filter on dose-response fit quality wherever it is reported, and estimate the between-protocol offset from compounds measured in common before considering any merge. The framework is sound and the code runs. Sample size is the limiting factor, closing the gap between 16 compounds and the 1911 used by published benchmarks is the work that remains, and until it is closed no ranking of grape pomace polyphenols produced this way should be relied on, including the one in Table 13.

Table 11
Published performance of eleven regression algorithms on 1911 DPPH IC_{50} records, as reported by Ghironi et al. [23]. Reproduced for benchmarking; no model in this table was trained here.

Algorithm	MAE	RMSE	Test R^2
Extra Trees	0.31	0.45	0.77
Gradient Boosting	0.32	0.46	0.76
XGBoost	0.33	0.47	0.75

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723177

Algorithm	MAE	RMSE	Test R^2
Random Forest	0.33	0.48	0.74
k-Nearest Neighbours	0.33	0.50	0.72
AdaBoost	0.47	0.61	0.58
Decision Tree	0.46	0.67	0.49
Ridge	0.51	0.70	0.44
Elastic Net	0.53	0.73	0.40
Bayesian Ridge	0.54	0.74	0.38
Lasso	0.56	0.78	0.31
Consensus of the top three	0.31	0.45	0.78

Errors are in pIC_{50} log units. These results used random splitting, so a scaffold-split result is not directly comparable.

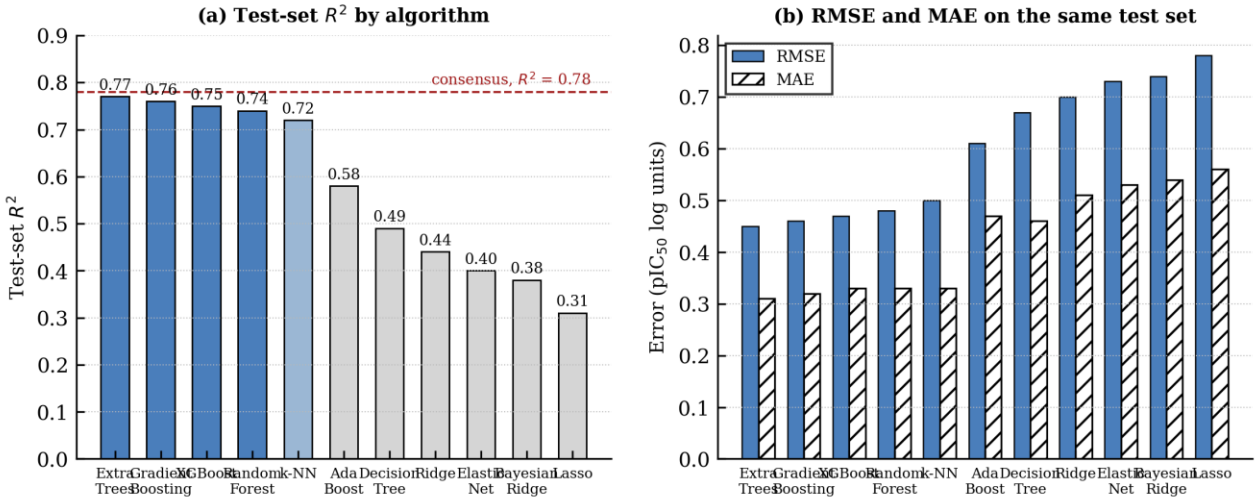

Fig. 12. Published DPPH regression benchmark on 1911 compounds, redrawn from the values of Ghironi et al. [23]. Ensemble trees beat linear models; the ceiling is about R^2 0.78.

Table 12
Published comparison of graph architectures on three antioxidant benchmarks, as reported by Cai et al. [30]. Values are classification accuracy from five-fold cross-validation.

Architecture	Parameters	AnOxPePred	AnOxPP	AOPP
GraphSAGE	23 k	0.7395	0.8392	0.8178
Graph attention network	54 k	0.7331	0.8160	0.7525
Graph isomorphism network	19 k	0.7082	0.6651	0.6452
Cheb Net	32 k	0.7357	0.8199	0.8131
Message passing neural network	75 k	0.7473	0.8703	0.8201

Data and code availability

The curated 16-compound corpus of Table 11, including canonical SMILES and generated InChIKeys, the measured literature values of Table 10, and the complete reference implementation of the curation protocol, molecular graph construction, message-passing network, baseline models and validation battery are supplied as supplementary material. The run reported in Section 7 is reproducible from that code with the seeds stated in Table 7. No trained model weights are distributed, because the model did not reach predictive performance on this corpus.

REFERENCES

[1] Tournour HH, Segundo MA, Magalhães LM, Barreiros L, Queiroz J, Cunha LM. Valorization of grape pomace: extraction of bioactive phenolics with antioxidant properties. *Ind Crops Prod* 2015;74:397-406. ScienceDirect

[2] Moutinho J, Gouvinhas I, Domínguez-Perles R, Barros A. Optimization of the extraction methodology of grape pomace polyphenols for food applications. *Molecules* 2023;28:3885. MDPI

[3] Zhu MT, Huang YS, Wang YL, Shi T, Zhang LL, Chen Y, et al. Comparison of (poly)phenolic compounds and antioxidant properties of pomace extracts from kiwi and grape juice. *Food Chem* 2019;271:425-32. ScienceDirect

[4] González-Centeno MR, Jourdes M, Femenia A, Simal S, Rosselló C, Teissedre PL. Characterization of polyphenols and antioxidant potential of white grape pomace byproducts (Vitis vinifera L.). *J Agric Food Chem* 2013;61:11579-87. ACS

[5] Rockenbach II, Rodrigues E, Gonzaga LV, Caliari V, Genovese MI, Gonçalves AESS, et al. Phenolic compounds content and antioxidant activity in pomace from selected red grapes (Vitis vinifera L. and Vitis labrusca L.) widely produced in Brazil. *Food Chem* 2011;127:174-9. ScienceDirect

[6] Rockenbach II, Gonzaga LV, Rizelio VM, Gonçalves AESS, Genovese MI, Fett R. Phenolic compounds and antioxidant activity of seed and skin extracts of red grape (Vitis vinifera and Vitis labrusca) pomace from Brazilian winemaking. *Food Res Int* 2011;44:897-901. ScienceDirect

[7] Pinelo M, Manzocco L, Nuñez MJ, Nicoli MC. Interaction among phenols in food fortification: negative synergism on antioxidant capacity. *J Agric Food Chem* 2004;52:1177-80. ACS

[8] Pinelo M, Rubilar M, Jerez M, Sineiro J, Núñez MJ. Effect of solvent, temperature, and solvent-to-solid ratio on the total phenolic content and antiradical activity of extracts from different components of grape pomace. *J Agric Food Chem* 2005;53:2111-7. ACS

[9] Ruberto G, Renda A, Daquino C, Amico V, Spatafora C, Tringali C, et al. Polyphenol constituents and antioxidant activity of grape pomace extracts from five Sicilian red grape cultivars. *Food Chem* 2007;100:203-10. ScienceDirect

[10] Negro C, Aprile A, Luvisi A, De Bellis L, Miceli A. Antioxidant activity and polyphenols characterization of four monovarietal grape pomaces from Salento (Apulia, Italy). *Antioxidants* 2021;10:1406. MDPI

[11] Radulescu C, Olteanu RL, Buruleanu CL, Nechifor (Tudorache) M, Dulama ID, Stirbescu RM, et al. Polyphenolic screening and the antioxidant activity of grape pomace extracts of Romanian white and red grape varieties. *Antioxidants* 2024;13:1133. MDPI

[12] Di Lorenzo C, Colombo F, Biella S, Orgiu F, Frigerio G, Regazzoni L, et al. Phenolic profile and antioxidant activity of different grape (Vitis vinifera L.) varieties. *BIO Web Conf* 2019;12:04005. BIO Web of Conferences

[13] Bors W, Michel C. Chemistry of the antioxidant effect of polyphenols. *Ann N Y Acad Sci* 2002;957:57-69. Wiley

[14] Baliyan S, Mukherjee R, Priyadarshini A, Vibhuti A, Gupta A, Pandey RP, et al. Determination of antioxidants by DPPH radical scavenging activity and quantitative phytochemical analysis of Ficus religiosa. *Molecules* 2022;27:1326. MDPI

[15] Blois MS. Antioxidant determinations by the use of a stable free radical. *Nature* 1958;181:1199-200. Nature

[16] Brand-Williams W, Cuvelier ME, Berset C. Use of a free radical method to evaluate antioxidant activity. *LWT Food Sci Technol* 1995;28:25-30. ScienceDirect

[17] Munteanu IG, Apetrei C. Analytical methods used in determining antioxidant activity: a review. *Int J Mol Sci* 2021;22:3380. MDPI

[18] Pérez-Jiménez J, Saura-Calixto F. Effect of solvent and certain food constituents on different antioxidant capacity assays. *Food Res Int* 2006;39:791-800. ScienceDirect

[19] Xie J, Schaich KM. Re-evaluation of the 2,2-diphenyl-1-picrylhydrazyl free radical (DPPH) assay for antioxidant activity. *J Agric Food Chem* 2014;62:4251-60. ACS

[20] Zhao D, Zhang Y, Chen Y, Li B, Zhou W, Wang L. Highly accurate and explainable predictions of small-molecule antioxidants for eight in vitro assays simultaneously through an alternating

multitask learning strategy. *J Chem Inf Model* 2024;64:9098-110. ACS

[21] Rice-Evans CA, Miller NJ, Paganga G. Structure-antioxidant activity relationships of flavonoids and phenolic acids. *Free Radic Biol Med* 1996;20:933-56. ScienceDirect

[22] Moriwaki H, Tian YS, Kawashita N, Takagi T. Mordred: a molecular descriptor calculator. *J Cheminform* 2018;10:4. Springer

[23] Ghironi S, Viganò EL, Selvestrel G, Benfenati E. QSAR models for predicting the antioxidant potential of chemical substances. *J Xenobiot* 2025;15:80. MDPI

[24] Kim HC, Ha SY, Yang JK. Modeling radical scavenging activity using molecular descriptors of unclassified compounds. *J King Saud Univ Sci* 2025;37:5552024. Journal of King Saud University - Science

[25] Goya Jorge E, Rayar AM, Barigye SJ, Jorge Rodríguez ME, Sylla-Iyarreta Veitía M. Development of an in silico model of DPPH free radical scavenging capacity: prediction of antioxidant activity of coumarin type compounds. *Int J Mol Sci* 2016;17:881. MDPI

[26] Muraro C, Polato M, Bortoli M, Aiolli F, Orian L. Radical scavenging activity of natural antioxidants and drugs: development of a combined machine learning and quantum chemistry protocol. *J Chem Phys* 2020;153:114117. AIP Publishing

[27] Ayres L, Benavidez T, Varillas A, Linton J, Whitehead DC, Garcia CD. Predicting antioxidant synergism via artificial intelligence and benchtop data. *J Agric Food Chem* 2023;71:15644-55. ACS

[28] Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural message passing for quantum chemistry. *Proc 34th Int Conf Mach Learn, PMLR* 2017;70:1263-72. arXiv

[29] Hakkal S, Ait Lahcen A. Leveraging graph neural network for learner performance prediction. *Expert Syst Appl* 2025;293:128724. ScienceDirect

[30] Cai J, Lou X, Chong CF, Alex D, Arrais JP, Wang Y, et al. Multi-AOP: a lightweight multi-view deep learning framework for antioxidant peptide discovery. *Bioresour Bioprocess* 2026;13:21. Springer

[31] Deng W, Chen Y, Sun X, Wang L. AODB: a comprehensive database for antioxidants including small molecules, peptides and proteins. *Food Chem* 2023;418:135992. ScienceDirect

[32] Weininger D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J Chem Inf Comput Sci* 1988;28:31-6. ACS

[33] Heller SR, McNaught A, Pletnev I, Stein S, Tchekhovskoi D. InChI, the IUPAC International Chemical Identifier. *J Cheminform* 2015;7:23. Springer

[34] Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. *Int Conf Learn Represent (ICLR)* 2017. arXiv

[35] Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y. Graph attention networks. *Int Conf Learn Represent (ICLR)* 2018. arXiv

[36] Xu K, Hu W, Leskovec J, Jegelka S. How powerful are graph neural networks? *Int Conf Learn Represent (ICLR)* 2019. arXiv

[37] Alexander DLJ, Tropsha A, Winkler DA. Beware of R-squared: simple, unambiguous assessment of the prediction accuracy of QSAR and QSPR models. *J Chem Inf Model* 2015;55:1316-22. ACS

[38] Bemis GW, Murcko MA. The properties of known drugs. 1. Molecular frameworks. *J Med Chem* 1996;39:2887-93. ACS

[39] OECD. *Guidance document on the validation of (quantitative) structure-activity relationship [(Q)SAR] models. OECD Series on Testing and Assessment No. 69.* Paris: OECD Publishing; 2007. OECD

[40] Landrum G. *RDKit: open-source cheminformatics software*; 2024. Available from: RDKit

[41] Iacopini P, Baldi M, Storchi P, Sebastiani L. Catechin, epicatechin, quercetin, rutin and resveratrol in red grape: content, in vitro antioxidant activity and interactions. *J Food Compos Anal* 2008;21:589-98. ScienceDirect

[42] Lv Q, Luo F, Zhao X, Liu Y, Hu G, Sun C, et al. Identification of proanthocyanidins from litchi (Litchi chinensis Sonn.) pulp by LC-ESI-Q-TOF-MS and their antioxidant activity. *PLoS One* 2015;10:e0120480. PLOS

[43] Luchian CE, Cotea VV, Vlase L, Toiu AM, Colibaba LC, Raschip IE, et al. Antioxidant and antimicrobial effects of grape pomace extracts.

BIO Web Conf 2019;15:04006. BIO Web of Conferences

[44] Sy B, Krisa S, Richard T, Courtois A. Resveratrol, ε-viniferin, and vitisin B from vine: comparison of their in vitro antioxidant activities and study of their interactions. *Molecules* 2023;28:7521. MDPI

[45] *Bioactive depside and anthocyanin compounds, compositions, and methods of use.* US Patent 8,846,757 B2; 2014. Available from: Google Patents

[46] Moazzen A, Öztinen N, Ak-Sakalli E, Koşar M. Structure-antiradical activity relationships of 25 natural antioxidant phenolic compounds from different classes. *Heliyon* 2022;8:e10467. ScienceDirect

[47] Xiong Z, Wang D, Liu X, Zhong F, Wan X, Li X, et al. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *J Med Chem* 2020;63:8749-60. ACS