A Convex Polygon Planarity

Here we present a proof of our planarity criterion for convex polygon drawings of hypergraphs (Theorem 2 from the main paper), which is inspired by Kuratowski's Theorem that states that a graph is planar if and only if it does not contain a subdivision of the complete graph K_5 or complete bipartite graph $K_{3,3}$ [34]. Similarly, Theorem 2 states that a hypergraph has a convex polygon representation if and only if it does not contain one of our forbidden sub-hypergraphs described in Section 4.2.1 (Figure 6) of the main paper.

Our proof requires three intermediate results: (1) the definition of a new graph representation corresponding to some polygon drawing of the hypergraph which we call the *face triangulation graph*, (2) a verification that the face triangulation graph meets the criteria of a convex representation provided by Thomassen [53], and (3) a proof that any articulation vertices in the hypergraph must appear on some face boundary of a convex polygon representation.

We first define connectedness and articulation vertices for hypergraphs as well as facial cycles in graph drawings. A graph is connected if there exists a path between every pair of distinct vertices. Connectedness for hypergraphs is defined similarly (see Bretto [10]). A graph is said to be *biconnected* if it does not contain any articulation vertices. An articulation vertex (also called cut-vertex) is a vertex whose removal makes the graph disconnected. We define biconnected hypergraphs and articulation vertices for hypergraphs in the same way. From here, it is natural to consider a biconnected component of a hypergraph: a maximal set of vertices $X \in V$ such that the sub-hypergraph induced by X is biconnected. Note that the graph consisting of a single edge and its two endpoint vertices is considered a biconnected graph. Similarly, we consider any hypergraph containing a single hyperedge and its incident vertices to be a biconnected hypergraph. Biconnected graphs and hypergraphs are central to our definition of the face triangulation graph and our first criterion for hypergraph convex polygon planarity. A *face* in a planar drawing of a graph is a region in the plane bounded by a set of vertices and edges. The unbounded region outside of the planar drawing is counted as the *exterior face*. The boundary of each interior face defines an *interior facial cycle*, and the exterior face defines the *exterior facial cycle*. We similarly define a face in a planar polygon drawing of a hypergraph as a region in the plane not covered by a hyperedge that is bounded by a set of vertices and polygon sides.

Our first claim regarding convex polygon representations requires a new definition for the *face triangulation graph* of a polygon drawing for a biconnected hypergraph. Let $H = \langle V, E \rangle$ be a biconnected, Zykov planar hypergraph. Then the König graph $K(H) = (X, Y, D)$ is a planar graph (Section 4.2.1). Recall that $K(H)$ is a bipartite graph containing a vertex $x \in X$ for each hypergraph vertex $v \in V$ and a vertex $y \in Y$ for each hyperedge $e \in E$ where $(x, y) \in D$ if the corresponding hypergraph elements v and e are incident in H. Let H have a polygon drawing determined by some planar representation of $K(H)$ where each vertex $v \in V$ has the same location in the plane as its corresponding vertex $x \in X$ (Figure 13 (a,b)). Let the vertices of each hyperedge polygon be ordered according to their angular coordinates relative to the corresponding vertex $y \in Y$. With this polygon drawing of H, the *face triangulation graph* $T(H)$ is constructed by the following procedure:

Procedure 3.

1. *Let $T(H)$ include all the vertices of H.*
2. *For two vertices $u, v \in V$, let (u, v) be an edge in $T(H)$ if (u, v) form the side of a polygon in the drawing of H.*
3. *Let $\{F_1, F_2, \ldots\}$ be the interior facial cycles in our current construction of $T(H)$ that correspond to a hypergraph face in the drawing of H (Figure 13 (b)).*
4. *For each interior facial cycle F_i, add a vertex c_i located in the interior of F_i and edges (c_i, v) for each vertex $v \in V(F_i)$ (Figure 13 (c)).*

With this definition, our goal is to show that the face triangulation graph has a planar drawing with convex facial cycles only for a specific class of hypergraphs. Thomassen [53] provides a characterization for such graph drawings which they term convex representations.

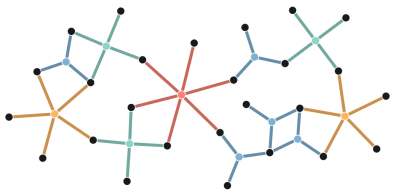

(a) Drawing of König graph $K(H)$.

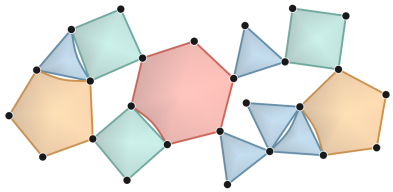

(b) Drawing of hypergraph H.

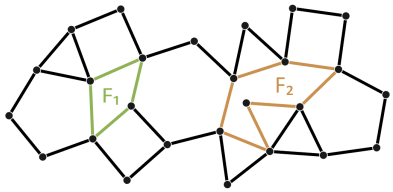

(c) Interior facial cycles $\{F_1, F_2\}$.

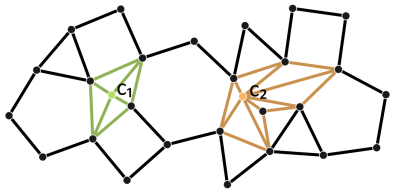

(d) Face triangulation graph $T(H)$.

Fig. 13: Construction of the face triangulation graph from a planar polygon drawing of a hypergraph.

Theorem 4 (Thomassen [53]). *Let G be a biconnected planar graph and let S be a cycle which is the face boundary of some plane representation of G. Let Σ be a convex polygon representing S. Then Σ can be extended into a convex representation of G if and only if*

(i) each vertex x in $G - V(S)$ of degree at least 3 is joined to S by three paths that are disjoint except for x,
(ii) each cycle which is edge-disjoint from S has at least three vertices of degree at least 3, and
(iii) no S-component has all its vertices of attachment on a path of S corresponding to a straight line segment of Σ.

For our polygon layouts of hypergraphs, we require that each hyperedge be drawn as a strictly convex polygon. Thus, we are interested in the case where the face triangulation graph has strictly facial cycles. Thomassen notes that if Σ is restricted to being strictly convex, condition (iii) becomes redundant. Thomassen also notes that every vertex $x \in V(G) \setminus V(S)$ with degree 2 must be on a straight line segment in any convex representation of G. It follows that the faces whose boundaries include x are not strictly convex. Thus, if we require that every face boundary be strictly convex, Thomassen's Theorem is reduced to the following:

Theorem 5 (Strictly Convex Representations of Graphs). *Let G be a biconnected planar graph and let S be a cycle which is the face boundary of some plane representation of G. Let Σ be a strictly convex polygon representing S. Then Σ can be extended into a strictly convex representation of G if and only if each vertex x in $G - V(S)$ is joined to S by three paths that are disjoint except for x.*

The proof of Theorem 5 follows from the proof of Thomassen's Theorem provided in [53]. These theorems on convex representations of graphs motivate an extension to convex polygon representations of hypergraphs. We use the face triangulation graph $T(H)$ to connect these theories on graph drawing to hypergraph polygon drawings. First, we must specify the conditions under which a face triangulation graph meets the prerequisites for Theorem 5.

Lemma 6. *If the polygon drawing of H corresponds to a plane representation of $K(H)$, then the face triangulation graph $T(H)$ is a planar graph.*

Proof. Since the König graph $K(H) = (X, Y, D)$ is planar, we know that it does not contain a subgraph homeomorphic to K_5 or K_{33}. Let us augment $K(H)$ by adding an edge (u, v) for every $u, v \in X$ such that (u, v) forms the side of a polygon drawing of H. This augmentation cannot create a subgraph homeomorphic to K_{33} since it does not add any bipartite edges. Notice that each hyperedge in H now corresponds to a wheel subgraph in $K(H)$ which is a planar graph. Since no edges are added between vertices in Y, it follows that a subgraph homeomorphic to K_5 cannot contain a vertex $y \in Y$ as a non-subdivision vertex. Further, each edge added between a pair of vertices $u, v \in X$ can already be obtained in $K(H)$ by smoothing their common adjacent vertex in Y. Thus, our augmentation of $K(H)$ does not affect its planarity. Now let us further augment $K(H)$ by removing all the vertices in X and all the edges in D. Then we are left with only the vertices corresponding to hypergraph vertices in H and edges corresponding to the sides of polygons in the drawing of H. Clearly, this does not affect the planarity of $K(H)$. We can now obtain the face triangulation graph $T(H)$ by triangulating each interior face of the augmented graph $K(H)$ that corresponds to a hypergraph face in the drawing of H. Since an interior face is a bounded region, it follows that we can place the new vertex inside the bounded region and add edges according to step 4 of Procedure 3 without introducing any edge crossings. Thus, our construction of $T(H)$ is a planar graph. □

Lemma 7. *If the polygon drawing of H corresponds to a plane representation of $K(H)$, then the face triangulation graph $T(H)$ is biconnected.*

Proof. Each hyperedge in $e \in E(H)$ is replaced by a cycle C_e in $T(H)$ following the sides of the corresponding polygon in the drawing of H. Each cycle C_e defines a biconnected subgraph of $T(H)$. Since H is biconnected by assumption, it follows that the union of all cycles C_e for $e \in E(H)$ is also biconnected. □

We now have the appropriate conditions to make the connection between Theorem 5 and convex polygon representations of hypergraphs.

Theorem 8. *Let H be a biconnected, Zykov planar hypergraph. Then H has a convex polygon representation if and only if it has a face triangulation graph $T(H)$ with a strictly convex representation.*

Proof. We first show that if H has a face triangulation graph with a strictly convex representation, then it has a convex polygon representation. Suppose there is a drawing of H such that the face triangulation graph $T(H)$ has a strictly convex representation. By the construction of $T(H)$, each hyperedge in H corresponds to a strictly convex facial cycle in the drawing of $T(H)$. It follows that if we draw each hyperedge in H as a polygon following the corresponding facial cycle in $T(H)$, we can obtain a convex polygon representation of H.

Now we show that if H has a convex polygon representation, then it has a face triangulation graph with a strictly convex representation. Suppose that H has a convex polygon representation. Let $T(H)$ be constructed from this representation of H according to Procedure 3. By Lemmas 6 and 7, we have that $T(H)$ is planar and biconnected. Let $T(H)$ be drawn according to the convex polygon representation of H, and let S be the exterior facial cycle of this drawing. Let Σ be a convex polygon representing S. Then by Theorem 5, Σ can be extended to a graph isomorphic to $T(H)$ with strictly convex facial cycles if and only if each vertex $x \in T(H) - V(S)$ is joined to S by three paths that are disjoint except for x. Let $x \in T(H) - V(S)$. Since $T(H)$ is biconnected, there must be at least two vertex-disjoint paths from x to $V(S)$. In the following paragraphs, we prove that x is joined to S through at least three disjoint paths by contradiction.

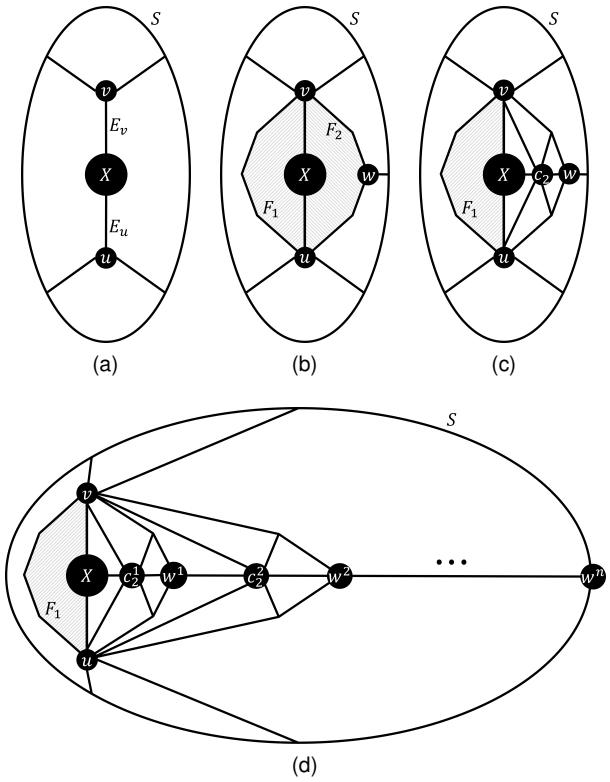

Fig. 14: Illustration of our argument for the proof of Theorem 8. In (a), we illustrate the subgraph X in the face triangulation graph which is joined to the exterior face S through two vertices u and v. In (b), we illustrate the faces F_1, F_2 enclosing X. In (c) we illustrate how the triangulation of F_2 connects X to a vertex w on F_2 through the triangulation vertex c_2. In (d) we illustrate how we construct a path from X to S through a sequence of face triangulations $(c_2^1, w^1, c_2^2, w^2, \ldots, c_2^n, w^n)$ that does not include u or v.

Suppose there are two vertices $u, v \in T(H)$ such that every path from x to $V(S)$ includes u or v. This implies that if u and v were removed, $T(H)$ would become disconnected, and x would be in a separate connected component from S. Let X be the connected component of $T(H)$ containing x when u and v are removed. Let E_u, E_v be the sets of edges between u and $V(X)$, and v and $V(X)$ respectively (Figure 14 (a)). With $T(H)$ drawn according to the convex polygon representation of H, it must be that X is drawn in the interior of S. It follows that X is enclosed by two faces F_1 and F_2 in the drawing of $T(H)$ whose boundaries contain u and v (Figure 14 (b)). Notice that F_1 and F_2 cannot both correspond to strictly convex polygons in the convex polygon representation of H since such polygons would necessarily share a side (u, v). This configuration would preclude X from being incident to both F_1 and F_2 while also being enclosed by F_1 and F_2. So, it must be that either F_1, F_2, or both are faces in the drawing of $T(H)$ that correspond to a hypergraph face in a convex polygon representation of H.

Consider the case where F_1 corresponds to a polygon in the convex polygon representation of H and F_2 does not. In order for F_2 to be drawn as a simple polygon, which must be the case since our drawing of $T(H)$ is planar, it must be that the boundary of F_2 contains at least one vertex $w \neq u, v$, $w \notin V(X)$ (Figure 14 (b)). Similarly, in the case where neither F_1 nor F_2 correspond to a polygon in the convex polygon representation of H, it must be that the boundary of either F_1 or F_2 contains at least one vertex $w \neq u, v$, $w \notin V(X)$. Without loss of generality, suppose that F_2 does not correspond to a polygon in the convex representation of H,

and that the boundary of F_2 contains such a vertex $w \neq u, v$, $w \notin V(X)$. In this case, the construction of $T(H)$ would have added a vertex c_2 to the interior of F_2 and edges (c_2, y) for every vertex y on the boundary of F_2. Then X would be connected to the vertex w through a path containing w which contradicts our observation that $w \notin X$ (Figure 14 (c)).

We can apply the same argument to an updated subgraph X and new faces F_1 and F_2 enclosing X. In this way, we can grow X with a sequence of vertices $(w_1, w_2, \ldots, w_n)$ until w_n is a vertex on S, which is possible assuming H and $T(H)$ are finite (Figure 14 (d)). Thus, we have shown that there must exist a path from x to $w_n \in V(S)$ that does not contain the vertices u or v. This contradicts our assumption that every path from x to S passes through u or v. Therefore, each vertex $x \in T(H) - V(S)$ is joined to S by three paths that are disjoint except for x, and Σ can be extended to a strictly convex representation of $T(H)$ by Theorem 5. □

We also wish to consider convex polygon planarity for hypergraphs that are not biconnected. To do this, we must address the placement of articulation vertices between biconnected hypergraph components, requiring the following lemma.

Lemma 9. *Let H be a hypergraph with exterior face S and interior faces $R = \{F_1, F_2, \ldots, F_n\}$ for some convex polygon representation of H. Then H also has convex polygon representations for each face $F_i \in R$ such that F_i is the exterior face and $R - F_i + S$ are the interior faces.*

This lemma can be proven in a similar manner to Theorem 8. Now we can extend Theorem 8 to a more general class of connected hypergraphs if we consider each biconnected component individually.

Theorem 10. *Let H be a Zykov planar hypergraph with k biconnected components. Let $\{B_1, B_2, \ldots, B_k\}$ be the sub-hypergraphs induced by the k biconnected components $V(B_i) \subseteq V(H)$. Then H has a convex polygon representation if and only if each sub-hypergraph B_i has a convex polygon representation where every vertex $x \in V(B_i)$ that is also an articulation vertex of H is located on a face boundary of some convex polygon representation of B_i.*

Proof. To prove Theorem 10 in the forward direction, suppose that each sub-hypergraph B_i has a convex polygon representation where every vertex $x \in V(B_i)$ that is also an articulation vertex of H is located on a face boundary of the convex polygon representation of B_i. Then we can construct a convex polygon representation for H by starting with the convex polygon representation of B_i. Now consider a sub-hypergraph B_j incident to B_i through the articulation vertex $x \in V(H)$. By Lemma 9, B_j has a convex polygon representation where the face boundary containing x is the exterior face boundary. It follows that we can draw B_j with this representation inside the face in B_i whose boundary contains x without introducing any polygon intersections (Figure 15). We can repeat this process until each of the biconnected sub-hypergraphs is drawn with an appropriate convex polygon representation.

To prove Theorem 10 in the reverse direction, suppose that H has a convex polygon representation. Clearly, a convex polygon representation of any sub-hypergraph B can be obtained by removing the vertices and hyperedges not in B from the convex polygon representation of H. Let x be an articulation in H belonging to biconnected sub-hypergraphs B_i and B_j. To reach a contradiction, suppose that x is not located on a face boundary of some convex polygon representation of B_i. Lemma 9 implies that x is not located on a face boundary for any convex polygon representation of B_i. It follows that if B_i is drawn with a convex polygon representation, there must be some intersection between a pair of hyperedge polygons $e_i \in E(B_i)$ and $e_j \in E(B_j)$ incident to the articulation vertex x. This contradicts our assumption that H has a convex polygon representation. □

Finally, we restate our main result for convex polygon representations of hypergraphs.

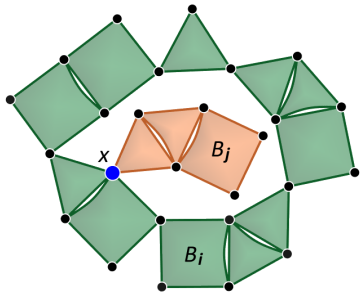

Fig. 15: Two biconnected sub-hypergraphs B_i and B_j incident through an articulation vertex x can be drawn without intersection if x is on a face boundary of both sub-hypergraphs.

Theorem 11. *(Theorem 2 from the main paper:) Let H be a Zykov planar hypergraph. Then H has a convex polygon representation if and only if it does not contain any of the following as a sub-hypergraph:*
(a) A 3-adjacent cluster of 2 hyperedges,
(b) A 2-adjacent cluster of 3 hyperedges,
(c) A strangled vertex,
(d) A strangled hyperedge.

Proof. $\Longrightarrow$ We prove the contrapositive statement: If H does not have a convex polygon representation, then it contains one of the forbidden sub-hypergraphs. By Theorem 10, H does not have a convex polygon representation if the sub-hypergraph induced by one of its biconnected components does not have a convex polygon representation. Let B represent such a biconnected sub-hypergraph. Then by Theorems 8 and 5, it must be that every face triangulation graph $T(B)$ drawn with exterior face boundary S contains a vertex $x \in T(B) - V(S)$ joined to S by fewer than three paths that are disjoint except for x. Without loss of generality, let $T(B)$ and S represent the drawing of the face triangulation graph of B containing the fewest such vertices x.

Consider the case where all of the face boundaries containing x represent hyperedges in B. If x is on exactly two such face boundaries, it must be that the corresponding hyperedges in B share at least 3 common vertices including x. This matches the definition of a 3-adjacent cluster of 2 hyperedges (Figure 16 (a)). If x is on more than two such boundaries, it follows that x is adjacent to at least three other vertices. At least one of these adjacent vertices, call it vertex y, must also be joined to S by fewer than three disjoint paths, otherwise, x would be joined to S by three disjoint paths. Without loss of generality, we can consider vertex y instead of vertex x, which may be contained in a different set of face boundaries, and could fall under one of the other following cases.

Now consider the case where x is on a face boundary that neither represents a hyperedge in H nor corresponds to a part of a hypergraph face in B. It follows that in the polygon drawing of B, x is positioned in the interior of some hyperedge polygon. This indicates the existence of a 2-adjacent hyperedge cluster of 3 hyperedges in B (Figure 16 (b)).

Now consider the case where x is on a face boundary corresponding to part of a hypergraph face in B. Then x is either one of the face triangulation vertices c from step 4 of Procedure 3, or is adjacent to such a vertex. If it is the latter, it follows that the adjacent face triangulation vertex c is also joined to S by fewer than three disjoint paths. Without loss of generality, assume that $x = c$. Then x is the central vertex of a wheel subgraph W in $T(B)$. Since x is adjacent to every other vertex in W, it follows that W is joined to the rest of $T(B)$ by fewer than three disjoint paths. This configuration corresponds to a strangled hyperedge in B (Figure 16 (c)).

Thus, we have accounted for all possible configurations of x, all of which indicate the existence of a forbidden sub-hypergraph in the biconnected sub-hypergraph B. Theorem 10 further implies that H does not have a convex polygon representation if it contains an articulation vertex x such that x does not appear on a face boundary of any convex polygon representation of some sub-hypergraph B induced by a biconnected component of H containing x. Then the hyperedges that

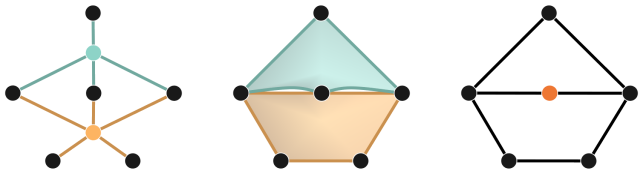

(a) 3-adjacent cluster of 2 hyperedges

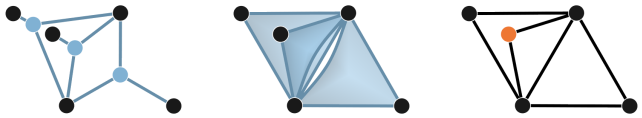

(b) 2-adjacent cluster of 3 hyperedges

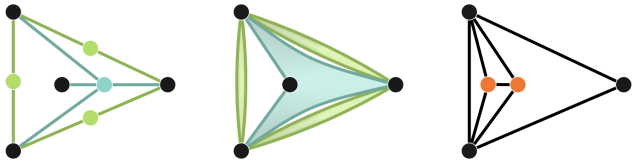

(c) Strangled hyperedge

Fig. 16: Forbidden sub-hypergraphs (middle) drawn according to a plane embedding of their König graphs (left) and their corresponding face triangulation graphs (right). The vertices highlighted in orange have fewer than 3 disjoint paths to the exterior face boundary.

are incident to x in B completely surround x in every convex polygon representation of B. This can only be possible if the hyperedges incident to x in B form a cycle in $B-x$. This matches the definition of a strangled vertex sub-hypergraph (Figure 6 (c) from the main paper).

$\Longleftarrow$ We prove the contrapositive statement: if H contains any of the forbidden sub-hypergraphs, it does not have a convex polygon representation. First, consider the case where H contains a 3-adjacent cluster of 2 hyperedges. When embedded in the plane, the three shared vertices in the 3-adjacent cluster must either form a triangle or be colinear. If they form a triangle, the intersection of two convex polygons containing the vertices must at least equal the area of the triangle. If the vertices are colinear, then the polygons containing them are not strictly convex. Thus, H does not have a convex polygon representation.

Now consider the case where H contains a 2-adjacent cluster of 3 hyperedges. Then the locations of the two shared vertices in the cluster define a line splitting the plane into two half-planes. Let $\{e_1, e_2, e_3\}$ be the three hyperedges in the cluster. Without loss of generality, e_1 can be drawn with its remaining vertices in one half plane, and e_2 can be drawn with its remaining vertices in the other half plane, and there is no intersection between the polygons for e_1 and e_2. For e_3 to be drawn as a convex polygon, its remaining vertices must be drawn in one half plane or the other, so it must have a nonzero intersection with the polygon for e_1 or e_2, and H does not have a convex polygon representation.

Now consider the case where H contains a strangled vertex $x \in V(H)$. Let C be the cycle among a proper subset of the vertices adjacent and hyperedges incident to x. If the vertices $V(C)$ are positioned in the plane such that their convex hull does not match their order in the cycle, it must be that the cycle crosses over itself and the drawing is non-planar. Otherwise, if x is located outside the convex hull of $V(C)$, it must be that one or more of the hyperedges in $E(C)$ have polygons crossing the interior and boundary of the convex hull, so the drawing is non-planar. If x is located inside the convex hull of $V(C)$, it follows that if each hyperedge in $E(C)$ is drawn as a convex polygon, the interior of the hull is completely tiled by these polygons. Thus, any other hyperedge polygon incident to x must intersect with one of the hyperedge polygons in $E(C)$, so H does not have a convex polygon representation.

Now consider the case where H contains a strangled hyperedge $e \in E(H)$. Let C be the cycle among a proper subset of the vertices incident and hyperedges adjacent to e. If the vertices $V(C)$ are located such that their convex hull does not match their order in the cycle, it must be that the cycle crosses over itself and the drawing is non-planar. Otherwise, if a vertex x incident to e but not in $V(C)$ is drawn inside the convex hull of $V(C)$, it follows that e cannot be drawn as a strictly convex polygon. If x is drawn outside the convex hull of $V(C)$, it follows that the drawing of e has nonzero intersection with at least one hyperedge polygon in $E(C)$, so H does not have a convex polygon representation. □

B Paper-Author Results

Figure 17: An enlarged version of Figure 10 from the main paper. A paper-author hypergraph dataset containing 1008 vertices and 429 hyperedges (a) is simplified with our framework down to the coarsest allowable scale H_{1214} (c) and the layout is optimized. Then the simplification is iteratively reversed, and the layout refined at each intermediate scale, an example of which is shown in (b), until the original scale H_0 is reached.

C Eye Tracking Gaze Paths

Figure 18: An enlarged version of Figure 11 from the main paper. Gaze fixation paths of two participants answering the same question for the trade agreement dataset in our user survey. The participant with the gaze path on the left did not study the dual hypergraph and answered the question incorrectly. The participant with the gaze path on the right studied both the primal and dual hypergraph visualizations and answered the question correctly.

D Eye Tracking Fixation Timelines

Figure 19: An enlarged version of Figure 12 from the main paper. Gaze fixation timelines of two participants answering a question in our user survey. The vertical axis indicates different regions on the participant's screen, including the question text and visualization scales. The horizontal axis represents the time in seconds that a participant spent on the question. The vertical lines in the plot indicate when the participant selected an answer. The blue lines indicate a correct answer and the red lines indicate an incorrect answer.

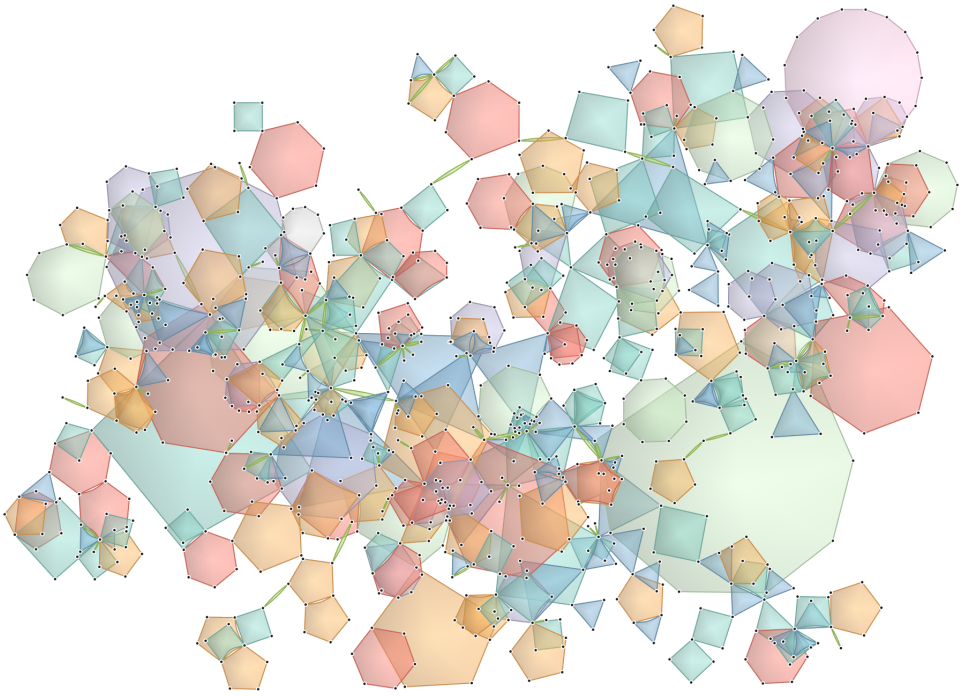

(a) Original Scale

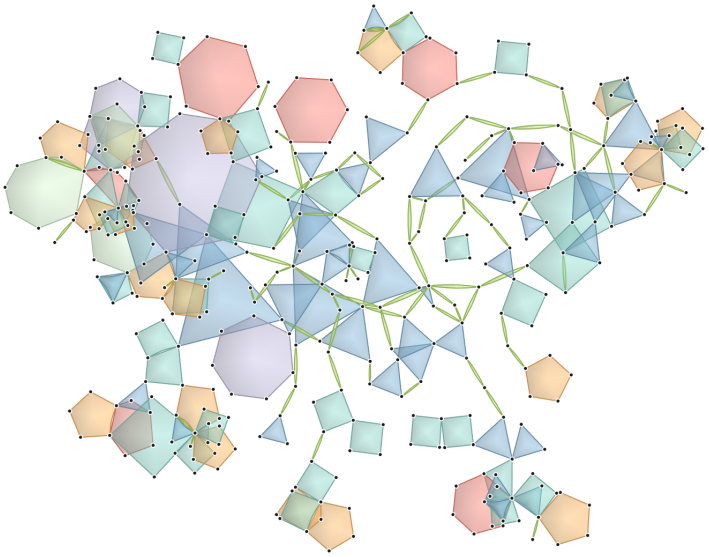

(b) Intermediate Simplified Scale

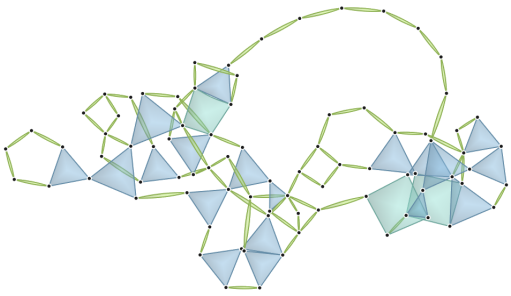

(c) Coarsest Simplified Scale

Fig. 17: Enlarged versions of the images in Figure 10 from the main paper. Final optimized layout of the original scale (a), coarsest simplified scale (c), and one intermediate simplified scale (b) for a paper-author hypergraph dataset.

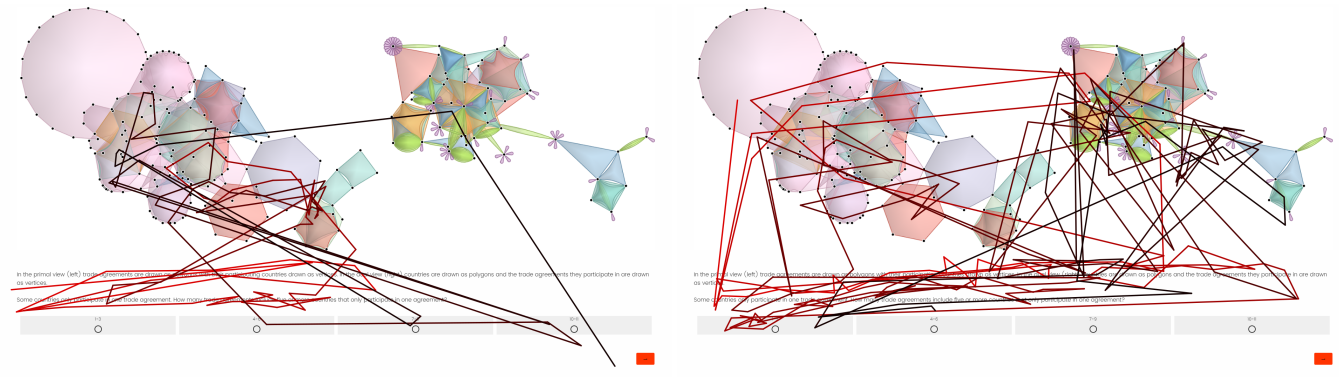

Fig. 18: Enlarged versions of the images in Figure 11 from the main paper. Gaze fixation paths of two user survey participants.

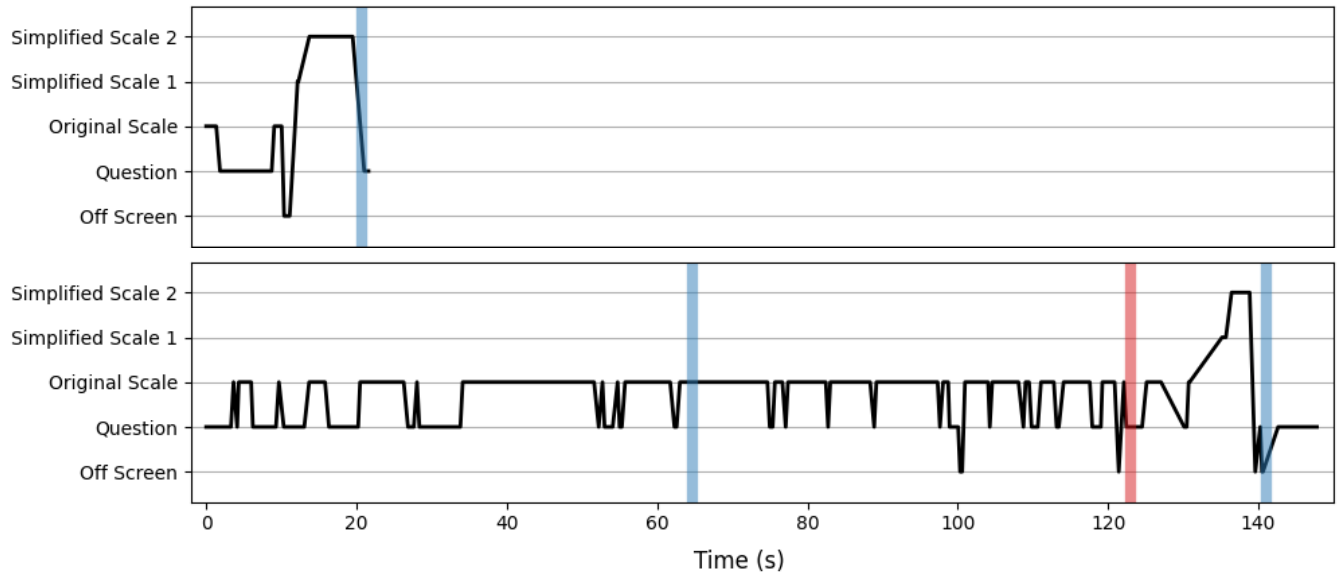

Fig. 19: Enlarged versions of the images in Figure 12 from the main paper. Gaze fixation timelines of two user survey participants.