Environmental Toxicology and Chemistry—Volume 37, Number 6—pp. 1734–1748, 2018
Received: 11 October 2017 | Revised: 11 December 2017 | Accepted: 24 February 2018

Environmental Toxicology

Adverse Outcome Pathway Networks II: Network Analytics

Daniel L. Villeneuve,[a,*] Michelle M. Angrish,[b] Marie C. Fortin,[c] Ioanna Katsiadaki,[d] Marc Leonard,[e] Luigi Margiotta-Casaluci,[f] Sharon Munn,[g] Jason M. O'Brien,[h] Nathan L. Pollesch,[a] L. Cody Smith,[i] Xiaowei Zhang,[j] and Dries Knapen[k]

[a]US Environmental Protection Agency, Mid-Continent Ecology Division, Duluth, Minnesota, USA
[b]US Environmental Protection Agency, National Center for Environmental Assessment, Research Triangle Park, North Carolina, USA
[c]Department of Pharmacology and Toxicology, Ernest Mario School of Pharmacy, Rutgers University, Piscataway, New Jersey, USA
[d]Centre for Environment, Fisheries and Aquaculture Science, Weymouth, United Kingdom
[e]L'Oréal Advanced Research, Aulnay-sous-Bois, France
[f]Institute of Environment, Health and Societies, Brunel University London, London, United Kingdom
[g]Joint Research Centre, European Commission, Ispra, Italy
[h]Environment and Climate Change Canada, National Wildlife Research Centre, Ottawa, Ontario, Canada
[i]Center for Environmental and Human Toxicology, University of Florida, Gainesville, Florida, USA
[j]State Key Laboratory of Pollution Control and Resource Reuse, School of the Environment, Nanjing University, Nanjing, People's Republic of China
[k]Zebrafishlab, Veterinary Physiology and Biochemistry, University of Antwerp, Wilrijk, Belgium

Abstract: Toxicological responses to stressors are more complex than the simple one-biological-perturbation to one-adverse-outcome model portrayed by individual adverse outcome pathways (AOPs). Consequently, the AOP framework was designed to facilitate de facto development of AOP networks that can aid in the understanding and prediction of pleiotropic and interactive effects more common to environmentally realistic, complex exposure scenarios. The present study introduces nascent concepts related to the qualitative analysis of AOP networks. First, graph theory–based approaches for identifying important topological features are illustrated using 2 example AOP networks derived from existing AOP descriptions. Second, considerations for identifying the most significant path(s) through an AOP network from either a biological or risk assessment perspective are described. Finally, approaches for identifying interactions among AOPs that may result in additive, synergistic, or antagonistic responses (or previously undefined emergent patterns of response) are introduced. Along with a companion article (part I), these concepts set the stage for the development of tools and case studies that will facilitate more rigorous analysis of AOP networks, and the utility of AOP network-based predictions, for use in research and regulatory decision-making. The present study addresses one of the major themes identified through a Society of Environmental Toxicology and Chemistry Horizon Scanning effort focused on advancing the AOP framework. *Environ Toxicol Chem* 2018;37:1734–1748. © 2018 The Authors. Environmental Toxicology and Chemistry published by Wiley Periodicals, Inc. on behalf of SETAC. This article is a US government work and, as such, is in the public domain in the United States of America.

Keywords: Adverse outcome pathway; Risk assessment; Predictive toxicology; Mixture toxicology; Adverse outcome pathway network; Network topology; Interactions

INTRODUCTION

The adverse outcome pathway (AOP) framework uses a modular structure to organize information concerning the linkage between a molecular-level perturbation of a biological system and the adverse outcome(s) that perturbation may cause. The modular structure is composed of 2 basic units: key events and key event relationships (Villeneuve et al. 2014a). These units are linked together in defined sequences (i.e., AOPs) that describe a series of measureable biological changes (key events) reflecting progression from a specific molecular initiating event to a defined adverse outcome, and the scientific evidence that supports the relationships (key event relationships) between those changes. Each AOP can be viewed as one biologically plausible, and scientifically defensible, chain of events that can lead from a defined biological perturbation to an adverse outcome. However, as more AOPs are described, systems or assemblages of interconnected AOPs that share one or more key events emerge. These assemblages are termed AOP networks (Knapen et al. 2018).

This article includes online-only Supplemental Data.
* Address correspondence to villeneuve.dan@epa.gov

Published online 28 February 2018 in Wiley Online Library (wileyonlinelibrary.com).
DOI: 10.1002/etc.4124

As part of a Society of Environmental Toxicology and Chemistry (SETAC)-sponsored Pellston Workshop that focused on advancing the AOP framework, the development and application of AOP networks was identified as one of the major themes that emerged from a review of more than 300 questions submitted by the scientific, risk assessment, and regulatory communities (LaLone et al. 2017). In the companion article (part I), Knapen et al. (2018) highlight the distinction between network-guided AOP development on the one hand (which follows the principles outlined for development of individual AOP descriptions), and deriving AOP networks from available information in the AOP knowledgebase on the other hand (Society for the Advancement of Adverse Outcome Pathways 2017). Data layers, analogous to those employed in geographic information systems, were proposed as a way to capture and represent the biological complexity underlying AOPs and influencing their dynamics (e.g., feedback/feedforward loops, modulating factors), without overly complicating the basic AOP framework. Filters and information layers derived from structured fields in the AOP knowledgebase were proposed as useful techniques for customizing the global AOP network. Finally, filtering and layering concepts were applied in a number of case studies to demonstrate how AOP networks can be built to provide answers to different types of questions (Knapen et al. 2018).

The present study builds on the concepts described by Knapen et al. (2018) and focuses more sharply on the issue of AOP network application in risk assessment, research, and decision-making. In particular, the study seeks to address how the network structure itself may be analyzed to derive information that can guide research and regulatory decision-making. First, we consider the similarities and distinctions between AOP networks and other types of complex networks that are commonly analyzed using techniques from graph theory (Trudeau 2013) and network science (Lewis 2009). Using 2 example AOP networks, we illustrate the application of a number of topology-based (Supplemental Data, Box 1) network analyses, and how they can guide our understanding of potential interactions among AOPs, as well as assay development and/or design of alternative testing strategies (e.g., integrated approaches to testing and assessment; Worth and Patlewicz 2016). We then discuss how different strategies may be used to identify the most important paths (called "critical paths"; Supplemental Data, Box 1) through AOP networks. Finally, we explore how AOP networks can provide first-order, qualitative insights into the potential interactions among AOPs, as well as the underlying ontological challenge that must be addressed to facilitate those analyses. More quantitative applications of AOP networks entailing incorporation of stressor-specific data and detailed consideration of response–response relationships defining transition from one key event to the next in an AOP network are not addressed in the present study. Together with the concepts and case studies presented by Knapen et al. (2018), these examples address prominent questions and themes (LaLone et al. 2017) concerning AOP networks and their general application.

AOPS AS NETWORKS

The term network broadly refers to any type of interconnected group. Many different types of networks can be represented graphically as systems of nodes (generally represented as dots or shapes) and edges (generally represented as lines; Table 1; Newman 2003; Lewis 2009; Trudeau 2013). The modular structure of the AOP framework (i.e., key events and key event relationships that can be reused in different AOPs) was developed with the construction of networks in mind (Villeneuve et al. 2014a; Organisation for Economic Co-operation and Development 2016a). In the case of AOP networks, key events are represented as nodes and key event relationships as directed edges (arrows).

Based on the mathematical study of graphs (graph theory; Pavlopoulos et al. 2011) and network science (Newman 2003; Caldarelli and Catanzaro 2012), common sets of principles and tools may be productively applied to analyze networks, including AOP networks. In fact, many network science–based tools have already become familiar to biologists and toxicologists: networks and network science have increasingly been used to represent and analyze molecular interactions or statistical associations among transcripts, genes, proteins, metabolites, and their regulatory factors. However, some important distinctions between AOP networks and other types of biological networks should be kept in mind (Table 2). First, in an AOP network, each node (key event) represents a measurable change (e.g., an increase or decrease) in the abundance of an object or state of a process compared with a reference, such as a control. As a result, a decrease in, for example, a hormone concentration, enzyme activity, or heart rate is represented as a separate key event from an increase in that hormone, enzyme activity, or heart rate. This has implications for how AOP networks are constructed and analyzed and how data may be layered on top of them. Second, most biological networks describe interactions at a single level of biological organization. By definition, AOP networks span multiple levels of biological organization. Finally, AOPs and AOP networks are focused on predictive utility over biological fidelity. The AOPs and the AOP networks are intended to help accurately predict how a system will respond to perturbation—even if that involves some degree of abstraction, simplification, or implicit embedding of more extensive biological understanding or structures within the key events and key event relationships represented. The objects (e.g., genes, proteins, cells, organs) included in an AOP description and measured as key events may often be the minimum set needed to support extrapolation or inference along the AOP. Detailed mechanistic understanding is only needed to the extent that it supports confidence in application. If the weight of correlative evidence is sufficient, an AOP can be collapsed to relatively few key events. As a result, considerable subjectivity is associated with the number of key event nodes that may be used to describe the path from a molecular initiating event to an adverse outcome.

To help illustrate the analyses and applications detailed in the present study, 2 AOP networks were derived from existing AOP descriptions within the AOP knowledgebase (Society for the Advancement of AOPs 2017; Supplemental Data, SI.1). For the first example, AOP 25 (aromatase inhibition leading to reproductive dysfunction; Villeneuve 2017) was used as a seed AOP, and an AOP network was derived to include all AOPs that shared at least one key event with AOP 25

TABLE 1: Examples of networks that are commonly represented graphically as systems of nodes and edges

Network type	Nodes	Edges
Transportation	Stations	Routes between stations
Computer	Computers and servers	Data transmission
Social	People	Relationships
Molecular biology	Genes and proteins	Interactions
Ecological/food web	Species	Energy flow
AOP	Key events	Key event relationships

AOP = adverse outcome pathway.

(Figure 1; hereafter termed the cytochrome P45019 [CYP19]-AOP network). Note that for the purposes of constructing the CYP19-AOP network, key events representing an opposite action (increase or decrease) on the same biological object or process were included. The second network (Figure 2; hereafter termed the thyroxine [T4]-AOP network) was generated by searching the AOP-Wiki (Society for the Advancement of AOPs 2017) for the term T4 (thyroxine, one of the thyroid hormones). The list of AOPs to include was compiled from the resulting AOP full text search results, and an AOP network was constructed from the list of key event relationships associated with those AOPs (Supplemental Data, SI.1). Although many of the topological analyses, identifications of critical paths, and considerations of interactions associated with these networks can be accomplished rather intuitively by visual inspection, these relatively simple examples help to demonstrate how various computational approaches could aid in the analysis of more complex AOP networks. They also help to highlight some of the challenges and limitations associated with the analysis of AOP networks.

Indeed, significant limitations need to be kept in mind whenever one is applying and interpreting AOP networks. First and foremost is the recognition that AOP networks are limited by the scope of knowledge and relationships currently captured in the AOP knowledgebase. The connections reflected in an AOP network are only those for which AOP descriptions have been developed. One cannot necessarily infer that a highly connected key event in an AOP network is more biologically important than a less connected key event. Such connections may simply reflect the fact that more extensive AOP development was centered around one particular key event than around others. Similarly, there is no objective way to define how many key events to include in a given AOP description (Villeneuve et al. 2014b). The level of abstraction or biological resolution to include in the AOP description is determined by the individual AOP developer, the scope of the available data, and the perception of the detail required to support an application. Thus, in applying the approaches discussed in the present study, it is important to remember that AOPs are not necessarily comprehensive representations of the biological system(s) that may be impacted by a given stressor or combination of stressors. Likewise, with the exception of AOPs that have gone through a formal review process by the Organisation for Economic Co-operation and Development or have been published in a peer-reviewed journal prior to (or in conjunction with) entry into the AOP-Wiki, there is no explicit quality assurance for AOPs in the AOP-Wiki. Instead, the AOP-Wiki relies on feedback from the scientific community for quality control (LaLone and Hecker 2017). Thus interpretations based on AOP networks need to be made with caution and with consideration of the quality and completeness of the underlying information.

AOP NETWORK ANALYTICS

Network topology

An AOP network effectively describes/represents a set of scientifically credible possibilities concerning 1) the diversity of biological perturbations that may cause a defined biological effect (e.g., an adverse outcome), and/or 2) the diversity of effects that may result from a stressor-induced biological change (e.g., a molecular initiating event). It also lays out the conditions

TABLE 2: Distinctive attributes of AOP networks compared with other types of biological networks (e.g., gene, protein, metabolite interaction networks; cell signaling networks)

Attribute	Typical biological networks[a]	AOP networks
Nodes	Discrete objects (e.g., a transcript, a protein, a cell, a species)	A measurable change in state (i.e., an increase or decrease) of an object or process (e.g., enzyme activity, circulation, behavior) compared with a reference state (e.g., a control)
Edges	Represent interactions at the same level of biological organization	Often involve transition from one biological level of organization to another
Intended fidelity	Representational focus—intended to accurately reflect the true/real structure of the system	Dynamical focus—intended to accurately represent and predict how a system will respond to perturbation, even if that involves some degree of abstraction, simplification, or implicit embedding of more detailed underlying systems understanding

[a]Typical of many biological networks; exceptions can be expected.
AOP = adverse outcome pathway.

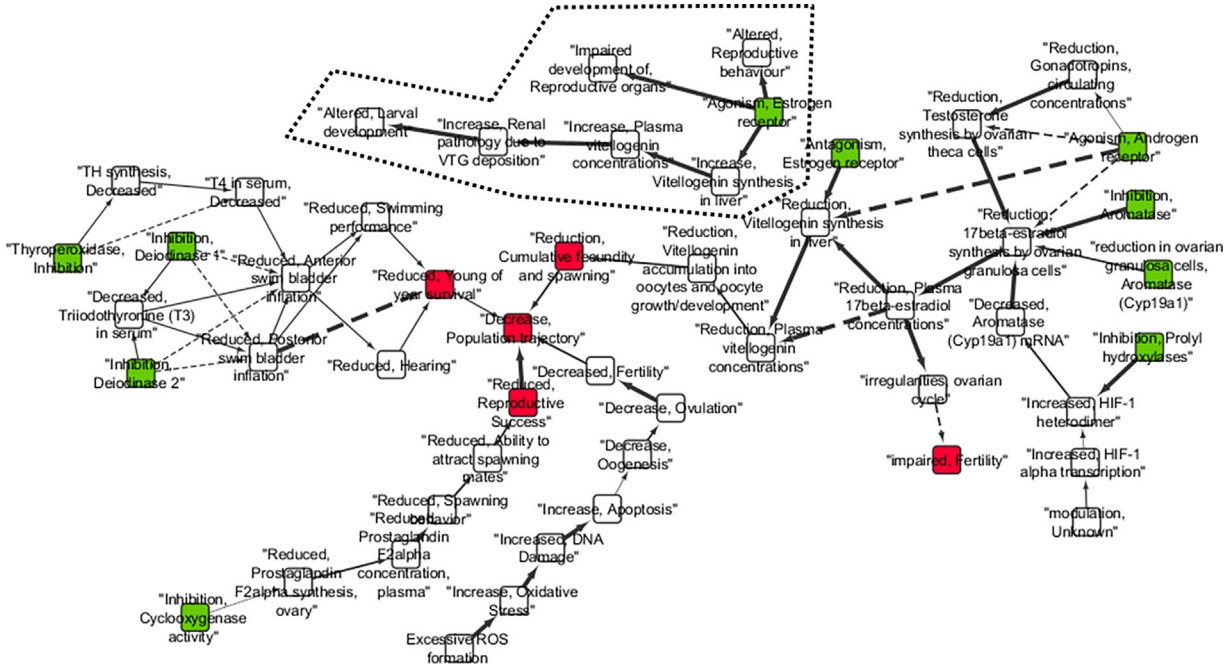

FIGURE 1: Example adverse outcome pathway (AOP) network 1 (cytochrome P45019 [CYP19]-AOP network). Shown is the network of all AOPs in the AOP-Wiki (Society for Advancement of AOPs 2017) that share at least one key event with those in AOP 25 (Villeneuve 2017). Rounded rectangles indicate key events. Arrows indicate key event relationships, with the arrow emanating from the upstream key event and into the downstream key event. Molecular initiating events are colored green. Adverse outcomes are colored red. Solid lines indicate relationships between key events that are adjacent in the sequence described in the AOP, and dashed lines indicate nonadjacent relationships. Arrow thickness indicates strength of evidence as defined in the AOP-Wiki for each key event relationship, where weak = thinnest arrows, strong = thickest arrows, and moderate = midsized arrows. A dotted line outlines a disconnected portion of the network. Unless noted otherwise, all key event titles and relationship information are directly as defined in the AOP-Wiki (Society for the Advancement of Adverse Outcome Pathways 2017; AOPs 25, 7, 23, 122, 123, 30, 29, 100, 155, 156, 157, 158, 159, and 216; Supplemental Data, Table S.1). The key event relationships shared by more than one AOP are shown as nonredundant (i.e., represented by a single arrow). VTG = vitellogenin; HIF-1 = hypoxia-inducible factor 1; TH = thyroid hormone; T4 = thyroxine; ROS = reactive oxygen species.

under which one possibility or another can be expected, noting that the various possibilities are not necessarily mutually exclusive. Although detailed explanations of the underlying support for those possibilities can be found in the descriptions of the individual AOPs, key events, and key event relationships (Organisation for Economic Co-operation and Development 2016a; Society for the Advancement of AOPs 2017), a great deal of information can be derived through examination and analysis of the overall structure of the network (i.e., the network topology). Conceptually, some of this information can be quickly gleaned from visual examination of the network graph, and it is recognized that an appropriate visualization and layout of a network is often a critical step in understanding (Newman 2003). However, as the networks become larger and more complex, it will be increasingly necessary to use various graph theory–based tools and computational algorithms to identify important and informative topological features. For computational purposes, AOP networks can be represented in a variety of forms such as adjacency matrices or from–to matrices (see the Supplemental Data). These can be readily processed using various software packages developed for network analyses (see Cytoscape 2017; the *graph*, *RBGL*, and *Rgraphviz* packages in Bioconductor, etc.). Custom tools specifically designed for analysis of AOP networks, such as the AOPXplorer (Burgoon 2017), are also under development.

Identifying points of convergence/divergence. Paths between key events within the AOP network can be described as either convergent, divergent, or mixed (Knapen et al. 2018). Identifications of convergent and divergent topologies are some of the most obviously useful analyses that one might apply to an AOP network. In a convergent topology (Supplemental Data, Box 1), key events from 2 or more AOPs are directed toward a common downstream key event (which could be the adverse outcome). Convergent motifs can help identify AOPs that may contribute in a joint manner toward the same outcome(s), suggesting potential additive or even synergistic effects if the upstream pathways are activated. In contrast, divergent structures branch off from a common key event toward a range of possible outcomes and can, for example, help define the pleiotropic effects a particular molecular initiating event or perturbation may have either within an individual organism or across different biological contexts (e.g., taxa, life stages, sexes, etc.). The key events that represent points of convergence in an AOP network (termed convergent key events; Knapen et al. 2018) could represent highly integrative endpoints that can detect the influences of a number of upstream perturbations. In contrast, a key event that serves as a point of divergence may be a measurement with particularly high predictive utility. Viewed from a drug development perspective, a drug that prevents a key event representing a point of convergence in an AOP network could potentially be designed to treat the effects of a

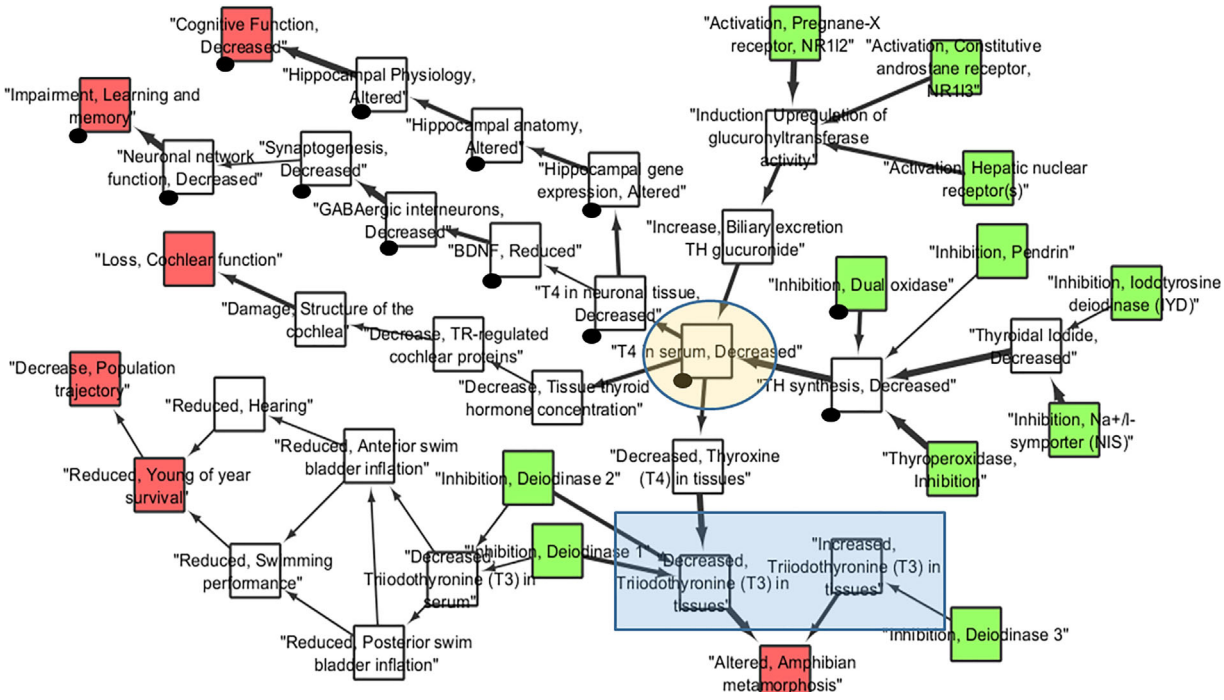

FIGURE 2: Example adverse outcome pathway (AOP) network 2 (thyroxine [T4]-AOP network). Shown is the network of 14 AOPs related to disruption of thyroid hormone signaling (Society for the Advancement of Adverse Outcome Pathways 2017; AOPs 8, 42, 54, 155, 156, 157, 158, 175, 188, 189, 190, 191, 192, and 193; Supplemental Data, Table S.2). Squares indicate key events. Arrows indicate key event relationships with the arrow emanating from the upstream key event and into the downstream key event. The key event relationships linking nonadjacent key events were filtered out of this network, and key event relationships shared by more than one AOP are shown as nonredundant (i.e., represented by a single arrow). In addition, to improve overall connectivity, the network was curated slightly with regard to titles and relationship information defined in the AOP-Wiki (see Supplemental Data, Table S.3 for details). Network overview is as follows. Molecular initiating events are colored green. Adverse outcomes are colored red. Arrow thickness indicates strength of evidence as defined in the AOP-Wiki for each key event relationship, where weak = thinnest arrows, strong = thickest arrows, and moderate = midsized arrows. A shaded circle highlights a key event that serves as the knot of a bow-tie motif within the network. A sequence of black dots highlights 2 examples of AOPs not described in the AOP-Wiki that emerge through network connectivity. A blue shaded rectangle highlights 2 key events that represent the same object, but different actions, within the AOP network. TH = thyroid hormone; GABAergic = γ-aminobutyric acid-ergic; BDNF = brain-derived neurotrophic factor; TR = thyroid hormone receptor.

range of upstream contributors to a disease state. Conversely, examination of AOPs that diverge from a key event that could be elicited by a drug treatment could help to identify potential side effects. Mixed structures may be characterized by points of convergence in the network that then diverge to a range of possible downstream outcomes. This divergence can often result in the "bow-tie" motif that is characteristic of many biological signaling networks, with the "knot" in the bow-tie often representing an important integrative biological signal (Friedlander et al. 2015). In the T4-AOP network, for example, the key event "T4 in serum, Decreased" (AOP-Wiki, Event 281; Society for the Advancement of AOPs 2017) represents the knot in a bow-tie motif (Figure 2). In the case of an AOP network, key events at the center of a bow-tie structure may represent a particularly important measurement that is frequently made in toxicological studies, an important control point in the biological system, or a biological change of particular interest to one or more AOP developers and potentially to risk assessors.

Computationally, points of convergence or divergence in an AOP network can be identified and quantified by various node ranking and centrality measures (Supplemental Data, Box 1; Huber et al. 2007; Pavlopoulos et al. 2011). For an AOP network, the "degree" (Table 3) of any key event is defined by the number of unique key event relationships it is linked to (Supplemental Data, Tables S.4 and S.5). For example, in the CYP19-AOP network, the key event "Reduction, 17beta-estradiol synthesis by ovarian granulosa cells" (AOP-Wiki Event 3; Society for the Advancement of AOPs 2017) is linked to 6 unique key event relationships (Figure 3A and Supplemental Data, Table S.4). In addition, because all the key event relationships in an AOP network are directed, the degree of each key event node in an AOP network can be further broken down in terms of the number of upstream ($degree_{in}$) or downstream ($degree_{out}$) key events it is linked to via key event relationships. The key event "Reduction, 17beta-estradiol synthesis by ovarian granulosa cells" is downstream of 5 key events, but upstream of just 1 key event (Figure 3A and Supplemental Data, Table S.4); therefore it would be a point of convergence in the AOP network. Key events linked to more upstream than downstream key events ($degree_{in} > degree_{out}$) can be broadly viewed as points of convergence in an AOP network, whereas those linked to more downstream than upstream key events ($degree_{out} > degree_{in}$) represent points of divergence.

It can also be informative to consider how many paths through the network pass through any given key event. In graph theory this is termed "betweenness" (Table 3 and Supplemental

TABLE 3: Overview of graph theory–based network analyses and their potential application(s) to AOP network analysis

Analysis/metric	Description	Potential use(s)
Node degree ($degree_{in}$ and $degree_{out}$)	The number of KERs linked to a KE. Directed networks, including AOP networks, can be broken down into $degree_{in}$ and $degree_{out}$, where $degree_{in}$ indicates the number of connections to upstream KEs, and $degree_{out}$ indicates the number of connections to downstream KEs.	Identify highly connected KEs within the overall AOP network. Identify points of convergence and/or divergence in the AOP network.
Betweenness centrality	Measure of the number of shortest paths between any KEs (*j*, k) in the AOP network that pass through the KE of interest (*i*) (Kitsak et al. 2007). Betweenness centrality is defined as: $g(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}}$ where $\sigma_{st}(v)$ is the number of shortest paths from *s* to *t* through KE *v* and σ_{st} is the number of shortest paths from *s* to *t*.	Identify important points of convergence/divergence in AOP networks. The KEs with high betweenness may represent measurements that are frequently made, critical control nodes within biological systems, or biological changes of particular interest to one or more groups involved in AOP development.
AOP Simple Path Occurrence	A variation on betweenness centrality that only considers the shortest path between MIEs and adverse outcomes, not between all pairs of KEs in the network. An AOP Simple Path Occurrence is defined as: $\Omega(v) = \sum_{(m \in M \neq v \neq a \in A)} \sigma_{ma}(v)$ where $\sigma_{ma}(v)$ is the number of shortest paths from *m* to *a* through node *v*, *M* is the set of MIEs, and *A* is the set of adverse outcomes. Note that, as for betweenness centrality (above), this calculation can also be normalized to the total number of shortest MIE-to-adverse outcome paths in the network. Normalized AOP Simple Path Occurrence: $\Omega(v) = \sum_{(m \in M \neq v \neq a \in A)} \frac{\sigma_{ma}(v)}{\sigma_{ma}}$	Identify important points of convergence/divergence in AOP networks. Identify KEs with high predictive value in terms of connecting many upstream MIEs to downstream AOs.
Topological sorting	A topological sort (or ordering) is a linear arrangement of the KEs of an AOP network such that for every KER {*KEi*,*KEj*} upstream *KEi* appears before downstream *KEj* in the arrangement (Skiena 1990; Weisstein 2017). This sorting can only be applied if there are no cycles in the network, such that KE_i could appear both before and after another KE_j.	Topological sorting can be useful for identifying the relative proximity of a KE in the network to the MIE (origin) or adverse outcome (terminus) of the AOPs it intersects with.
Eccentricity	The maximum shortest path length between KE_i and another KE_j in the network, defined as: $ecc = \max\{dist\ (i,j)\}$ Note that eccentricity is distinct from eccentricity centrality (*Cecc*), which is the inverse: $Cecc = \frac{1}{\max\{dist\ (i,j)\}}$ (Pavlopoulos et al. 2011; Cytoscape 2017; Netzwerkin 2017).	Useful for identifying the most upstream and downstream KEs in the network. Can also help identify upstream KEs that are connected to AOPs for which larger numbers of KEs have been defined.
Contraction	Identifies and removes cycles from the network, generating a graph that can be analyzed using algorithms applicable only to acyclic graphs.	Identification of cycles in an AOP network that may represent important features such as positive or negative feedback loops or modulating factors intrinsic to the AOP network.
Connectivity	A metric that indicates the number of connected components relative to the overall size of the network, where *E* = the number of KERs and *N* = the number of KEs. $C = \frac{E}{N(N-1)}$	The overall connectivity within an AOP network may indicate the relative potential for known complex interactions, and thus the potential uncertainty in predicting outcomes along those networks.
Matching index	A measure of how similar 2 KEs are within an AOP network based on the number of common neighbors they share, defined as: $M_{ij} = \frac{\Sigma nc}{\Sigma nt}$ where *nc* = neighbors common to KE*i* and KE*j*, and *nt* = total neighbors of KE*i* and KE*j* (Pavlopoulos et al. 2011).	Can be used to cluster KEs that are connected to similar upstream and/or downstream biology within an AOP network and sort KEs based on similarities in toxicological function.

AOP = adverse outcome pathway; KE = key event; KER = key event relationship; MIE = molecular initiating event.

Data, Figure S.5 and Tables S.5 and S.6). Traditionally, betweenness calculations consider the number of shortest paths (between every pair of nodes in the entire network) that have to pass through the node in question. However, in the context of an AOP network, the paths between molecular initiating events and adverse outcomes, rather than just any 2 key events in the network, are of particular interest. Consequently, using a custom R code (AOP Simple Path Occurrence; Table 3), we derived a variation on the betweenness calculation that considers key events specifically with regard to the number of paths between a

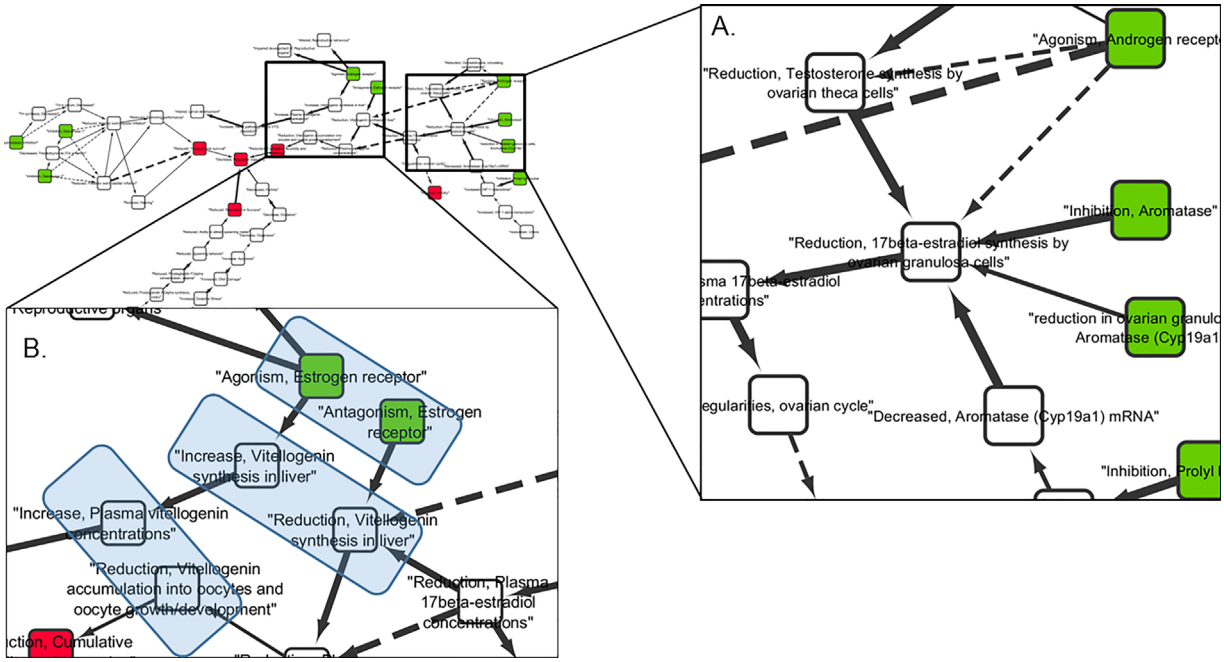

FIGURE 3: Adverse outcome pathway (AOP) network example 1 (cytochrome P45019 [CYP19]-AOP network) with view zoomed in on key features. (**A**) Zoomed-in view illustrating degree for the key event titled "Reduction, 17beta-estradiol synthesis by ovarian granulosa cells" (AOP-Wiki, Event 3; Society for the Advancement of AOPs 2017). (**B**) Zoomed-in view of several pairs of key events, highlighted by blue boxes, that represent the same object, but different actions, within the AOP network.

molecular initiating event and an adverse outcome that pass through them. Based on this calculation, we can see that the T4-AOP network (Figure 2) "T4 in serum, Decreased" (AOP-Wiki Event 281; Society for the Advancement of AOPs 2017) has the greatest AOP Simple Path Occurrence (Supplemental Data, Figure S.6), whereas in the main connected portion of the CYP19-AOP network (Figure 1), "Decreased, population trajectory" (AOP-Wiki Event 360; Society for the Advancement of AOPs 2017) has the greatest AOP Simple Path Occurrence (Supplemental Data, Figure S.7). Key events with high AOP Simple Path Occurrence are likely to be "hubs" in the overall AOP network and thus may be particularly useful to measure or to manipulate experimentally.

Eccentricity and topological sorting. The intent of topological sorting (Table 3) is to order the nodes in a directed network such that earlier nodes in a sequence are listed or displayed before later nodes when the nodes are arranged in a line. Individual AOPs are inherently sorted this way based on the sequence of cause–effect relationships they outline. Consequently, for a single AOP or a small AOP network, the causal sequence of key events is easy to discern based on the key event relationships (arrows; see Figures 1 and 2). However, as the AOP networks grow in size, even when the molecular initiating events and adverse outcomes are highlighted (e.g., color in Figures 1 and 2), it can become increasingly difficult to identify where various key events lie in terms of relative position along the sequence(s) from a molecular initiating event to an adverse outcome. Likewise, it can be difficult to pick out which key events even lie a long the paths that can trigger a given adverse outcome. Nevertheless, a number of computational approaches can help to reveal the ordering and relationships among key events represented in the network.

Calculation of key event "eccentricity" (Table 3) is one approach that can be used to discriminate molecular initiating events, or the furthest upstream key events from adverse outcomes, or the furthest downstream key events (Supplemental Data, Figure S.8 and Tables S.8 and S.9). Because AOP networks are directed networks, the farther downstream a key event is, the greater its eccentricity score will be (and the lower the inverse, eccentricity centrality, will be), because it is increasingly difficult to draw a directional path from that key event to another key event in the network. However, because calculation of eccentricity depends on path length (a somewhat subjective result of the number of key events an AOP developer includes in an AOP description), the results can be misleading if there are AOPs with a wide range of different lengths (numbers of key events included) in the network. For example, in the T4-AOP network, even though "Inhibition, Deiodinase 3" (AOP-Wiki, Event 1153; Society for the Advancement of AOPs 2017) is a molecular initiating event, it has a lower eccentricity score than many other key events in the network because one can only draw a path to 2 other key events in the network before reaching a terminal adverse outcome (Supplemental Data, Figure S.8).

Topological sorting (Skiena 1990; Weisstein 2017; Table 3) is an alternative approach toward ordering key events based on their causal and dependent relationships that is less impacted by subjectively defined path lengths (Figure 4 and Supplemental Data, Figures S.9 and S.10). Although topological sorting can

yield a solution that is nonunique, the dependent order of key events in the network is always maintained such that no causal (upstream) key events are positioned before their dependent (downstream) key events. Thus, for an AOP network that has been topologically sorted, if the analyst chooses a single key event, that key event can depend, at most, on the set of key events that precede it in the sorting. It cannot be dependent on any key events positioned further downstream.

One important caveat to topological sorting is that it can only be applied to acyclic graphs (e.g., no cycles such as feedforward loops). Cycles may occur in some AOP networks, even if there are no cycles in the individual AOPs from which the network was derived. The most common cause of cycles in an AOP network is the presence of a key event$_i$ that is upstream of key event$_j$ in one AOP, but downstream of key event$_j$ in another. This situation effectively introduces a 2-way arrow into the network, even though each of the original AOPs was unidirectional. To perform topological sorting, the cycles must be removed and represented as a single node on the graph through a process termed "contraction" (Table 3 and Figure 4). Given that contraction of 2 or more key events involved in a cycle results in a single node in the graphical depiction of the network, protocols for assigning attributes (e.g., title, shape, color, etc.) to the resulting contracted key events need to be defined so that they are apparent and interpretable. Nonetheless, the contraction process itself can be informative, because it provides a rapid, computational means to identify potential feedback or cross-talk processes that are intrinsic to an AOP network and may have an important bearing on the ability to predict outcomes. It is notable that numerous questions submitted to the SETAC Horizon Scanning exercise were focused on methods to identify and describe feedback and feedforward processes using the AOP framework (LaLone et al. 2017).

Connectivity. Evaluation of overall connectivity (Table 3) in AOP networks may also have some utility. Connectivity measures evaluate the relative extent to which there is a directed path from one key event to any other key event in the network. The AOP networks with greater connectivity indicate greater potential for known complex interactions, and thus greater potential uncertainty in predicting outcomes along those networks. In contrast, lower connectivity may suggest that there are just a few points of interaction that may need to be considered when one is inferring effects along the AOP network.

Matching indices, clustering, and network motifs. Particularly as the number of AOPs in the AOP knowledgebase grows and the associated AOP network expands, analyses that computationally identify AOPs sharing many of the same key events may also be quite useful. A matching index (Pavlopoulos et al. 2011; Table 3) can be applied to sort key events with regard to their similarity to one another. Likewise, a variety of common clustering approaches (e.g., neighbor joining) can be applied to adjacency matrices or other computational representations of AOP networks to identify key events or AOPs that share many of the same links and then cluster them away from those that are very independent (Pavlopoulos et al. 2011). Thus, for example, if one was observing effects associated with a particular key event or a common set of key events, clustering could be applied to identify the group of AOPs one might want to consider in an assessment.

Finally, there are computational approaches that can be used to highlight the presence of recurring patterns, or network motifs, that appear in a network significantly more than in a randomized network (Milo et al. 2002). Common motifs found in other types of molecular biological networks include feedforward and feedback loops, diamond structures, and even more complex structures (Vazquez et al. 2004; Alon 2007). However, in many cases, features like feedback loops or signal transduction cascades are not directly represented as key events in an AOP, but rather are embedded in the description of the biological plausibility and/or quantitative understanding of a key event relationship. Thus, it is unclear what types of over-represented motifs may be discovered as AOP networks are analyzed. This should prove to be an interesting research area in upcoming years.

Key event adjacency and topology-based analyses. In the AOP-Wiki, "nonadjacent key event relationships" (Supplemental Data, Box 1), previously termed "indirect key event relationships" (Organisation for Economic Co-operation and Development 2016a), are often created as a means to capture evidence that may skip over one or more of the key events in the

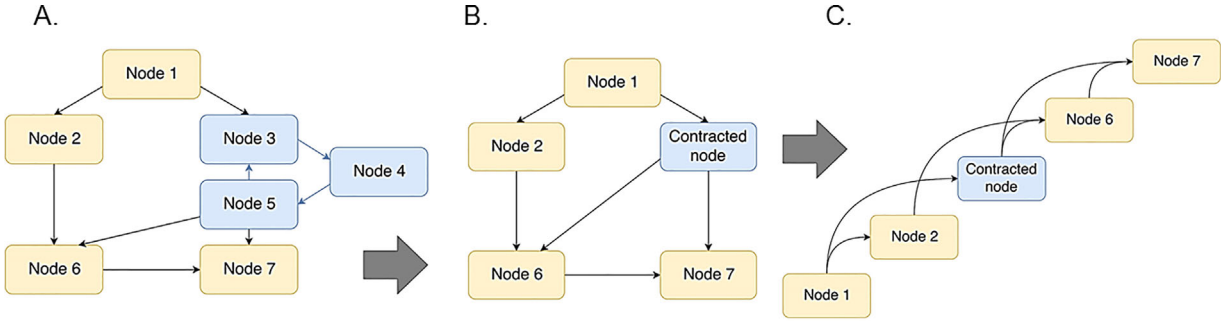

FIGURE 4: Generic example illustrating contraction and topological sorting of a network. (**A**) Generic directed network graph containing a cycle (key events 3, 4, and 5). (**B**) Graph of the same network following contraction of key events 3, 4, and 5, into a single contracted key event. Contraction results in a directed acyclic graph. (**C**) Graph of the contracted adverse outcome pathway network following topological sorting.

pathway. The creation of key event relationship descriptions for key events that are not next to one another in the sequence defined for an AOP helps to more fully capture the weight of evidence supporting an AOP while maintaining the modular structure of the knowledgebase and framework (Organisation for Economic Co-operation and Development 2016a). However, these nonadjacent key event relationships create some challenges in using topology-based analyses. Specifically, inclusion of nonadjacent key event relationships in an AOP network can inflate node degree and betweenness centrality or deflate distance-based calculations like eccentricity (Table 3) by introducing artificially short paths into the network (for examples, see Supplemental Data, Tables S.4–S.9). One solution is to filter the network to only include adjacent key event relationships when calculating topology-based analyses. However, such filtering can be challenging in the sense that adjacency is AOP specific. Because there is subjectivity in the number of key events included in an AOP description, a pair of key events that are adjacent in one AOP may in be nonadjacent in another. In general, if a key event relationship connects adjacent key events in any AOP in the knowledgebase, it should be viewed as an "adjacent key event relationship" for the purposes of network derivation. A comparison of results for the filtered versus unfiltered network (excluding versus including nonadjacent key event relationships) may be informative for certain questions. For example, in cases in which the number of key event relationships connected to a key event increases significantly when nonadjacent key event relationships are included, the results could be indicative of key events for which greater amounts of empirical evidence are likely available (i.e., key events that have been measured more frequently than other key events in the network).

Critical paths

By providing a framework for the description of the overall landscape of potential adverse outcomes resulting from particular biological perturbations, AOP network analyses can enable strategic identification of paths that have the greatest biological likelihood and/or relevance for risk assessment. This can in turn aid in the identification of endpoints with good predictive value that can serve as useful alternatives to the direct measurement of apical adverse outcomes (Organisation for Economic Co-operation and Development 2016b). For the purposes of our discussion, the path through an AOP network that is considered to be most important to an assessment or research question (and/or most dominant or biologically significant) was termed the "critical path" (Supplemental Data, Box 1). In the present study we distinguish "path" from "pathway" to recognize that the critical path from a key event to an adverse outcome within the AOP network may not necessarily follow the sequence laid out in an individual AOP in the knowledgebase, nor does it necessarily equate to a defined biological pathway, but rather may be a path that emerges only through the assembly and consideration of the interactions between multiple AOPs (Figure 5A). For example, in the AOP-Wiki, no AOPs linking inhibition of dual oxidases to decreased cognitive function or impaired learning and memory in mammals have been described to date. However, within the T4-AOP network (Figure 2), these emerge as potential paths through which an effect could occur.

Perspectives on what constitutes the critical path of interest can vary widely based on the regulatory context and/or mandate under which a risk assessment might be done, the research question an investigator may be interested in, or the type of application for which an AOP network is used. As a result, there is no one-size-fits-all approach to critical path identification. However, some of the more common types of data or information that may drive a critical path analysis are outlined in the following sections.

Problem formulation–defined critical paths. Problem formulation is the first step in environmental risk assessment; the scope and goals of the assessment are defined, a clear articulation of the question to be addressed is developed, and measurement endpoints are defined (US Environmental Protection Agency 1992, 1998). In many cases, up-front problem formulation (US Environmental Protection Agency 1992, 2014) can significantly reduce the range of paths that one would consider within a given AOP network, thereby aiding identification of the critical paths. Many of the filters and layers discussed by Knapen et al. (2018) can be effectively applied to narrow the range of possibilities and derive an AOP network that is fit-for-purpose. This allows for derivation of an AOP network that retains pathways relevant to the scope, question, and measurement endpoints defined while eliminating less relevant or ancillary pathways from the network so as not to overcomplicate the analysis. For example, if one were charged with conducting a risk assessment focused on cancer risk in humans, the network could be filtered to display only those paths that are relevant to humans (based on taxonomic applicability annotations associated with the key events and key event relationships) and linked via one or more paths through the network to the outcome of increased tumor formation.

The relevance of the apical adverse outcome in different regulatory contexts can be a particularly important factor to consider. Although all AOPs will ostensibly extend to an outcome that is accepted as adverse, this may not be true for all possible risk contexts and purposes. For example, an outcome like liver fibrosis may be of concern for a human health risk assessment, but would not necessarily be considered in an ecological risk assessment if a strong link to survival, growth, or reproduction was not established. Likewise, whereas a researcher may be interested in links between a molecular initiating event and behavioral effects that are plausibly related to impaired survival or reproduction, a risk assessor may be unable to consider such data unless the relevance to a population-level assessment can be established by strong empirical support. In essence, up-front problem formulation and scoping are needed to define the fit-for-purpose critical path(s) in an AOP network for a given application.

Weight-of-evidence–defined critical paths. Critical paths may also be identified based on the weight of evidence supporting the key event relationships between the key events.

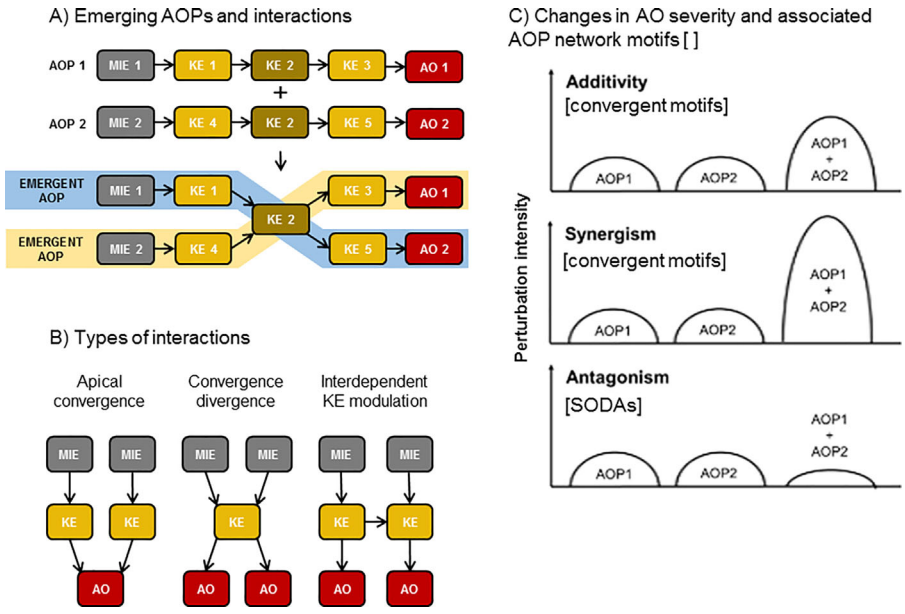

FIGURE 5: Generic illustration of various types of interactions relevant to the analysis of adverse outcome pathway (AOP) networks. (**A**) Graphical depiction of emergent AOPs that can arise when individual AOP descriptions are linked as an AOP network. (**B**) Illustration of some common types of interactions found in AOP networks. (**C**) Illustration of how AOP interactions may impact the intensity of perturbation of the key events (KEs) downstream of the point of interaction, and AOP network motifs that would commonly be associated with those interactions. SODA = same object, different action, where object and action are ontology terms that are used in defining a key event; MIE = molecular initiating event.

As AOPs are being described in the AOP knowledgebase, weight-of-evidence or confidence calls are being made by the AOP developer(s) (Becker et al. 2015; Organisation for Economic Co-operation and Development 2016a). These calls of high, medium, or low confidence in a key event relationship and high, medium, or low understanding of the quantitative nature of the relationship between each pair of key events may help to identify the critical paths that a risk assessor might be interested in. For example, if an authority wants to use an in vitro bioassay result as the basis for a hazard assessment in a regulatory setting, that authority may require a relatively high confidence in that measurement as a scientifically credible and defensible predictor of likelihood to cause the adverse outcome(s) (e.g., all the weight-of-evidence calls are ranked high or at least moderate). The assessor could prioritize the critical paths for digging into the details of the evidence supporting the key event relationships and ensuring that the data are of adequate quality to support extrapolation of the in vitro data to probable apical hazard(s), ignoring the paths for which only weak evidence has been assembled. In contrast, a researcher may be specifically interested in those parts of the AOP network where evidence is weakest, which may represent critical data gaps that could be important to address.

To aid in this process, a key event relationship confidence assessment filter or layer (Knapen et al. 2018) could be applied to an AOP network (forming a weighted network) prior to analysis, to ensure that only high-confidence paths are considered. For example, in AOP network 2 (Figure 2), the thickness of the key event relationships (arrows) reflects whether low, medium, or high confidence was assigned to each, with high confidence represented by the thickest arrows. Based on visual inspection of the network, the paths linking sodium iodide symporter (AOP-Wiki, Event 424; Society for the Advancement of AOPs 2017), or thyroperoxidase inhibition (AOP-Wiki, Event 279; Society for the Advancement of AOPs 2017) to decreased cognitive function (AOP-Wiki, Event 402; Society for the Advancement of AOPs 2017), or altered amphibian metamorphosis (AOP-Wiki, Event 1101; Society for the Advancement of AOPs 2017) have higher average key event relationship confidence calls compared with the path linking iodotyrosine deiodinase inhibition (AOP-Wiki, Event 1152; Society for the Advancement of AOPs 2017) to the same adverse outcomes (Supplemental Data, Figure S.11). Thus, an assessor may be comfortable using thyroperoxidase and sodium iodide symporter assay data, but not iodotyrosine deiodinase assay data, as a basis for hazard identification. Conceptually, this type of weight-of-evidence analysis could be done computationally, but decisions regarding how to weight adjacent and nonadjacent key event relationships along a path and how best to sum, average, or normalize the weight of evidence along each path would need to be defined.

In a risk assessment, all available studies are reviewed, and the quality of the data for each study is taken into consideration in an effort to identify critical effect and the best point of departure. When this process is translated to an AOP network, easily measured key events that link to paths for which all downstream edges indicate high confidence and/or high quantitative understanding may lend themselves to prioritization of sentinel, measurable, or critical endpoints. Depending on the needs of the assessment, if additional confidence were needed, the downstream key events along the path point to confirmatory endpoints to include in the assessment.

Similarly, for a given risk assessment, the assessor may have a defined set of available data. These data may align with a range of different key events in the AOP network, and some of the data may be more reliable (and thus weighted more heavily) than others. Because it would generally be preferred to weight the risk assessment more heavily toward the outcomes that align with higher confidence data, analysis of the AOP network (e.g., matching indices or clustering) could guide determination of the paths linked to high-confidence data, which in turn could define the critical path(s) to consider in the assessment. In principle, there is no reason that multiple criteria (i.e., confidence in the key event relationships, confidence in the data aligning with various key events that lead to the critical effect, and taxonomic relevance to the problem formulation) could not be considered together, allowing for even further refinement and definition of the most critical path for an assessment.

Biologically or toxicologically defined critical paths. Most of the considerations for defining critical paths described in the Problem formulation–defined critical paths and Weight-of-evidence–defined critical paths sections relate to the goals and objectives of certain steps in the risk assessment process (e.g., problem formulation or hazard identification). However, a number of considerations are more intrinsic to the biology that may define or further refine critical paths. For example, for a stressor that can trigger multiple pathways (e.g., polypharmacology), the critical path may be defined based on the molecular initiating event triggered at the lowest dose (i.e., defined by the relative sensitivity of the different pathways). Pathways that are only triggered at concentrations 10 or 100 times greater than others could plausibly be much less relevant to consider under most exposure scenarios. Generally speaking, pathways that are more readily perturbed by a given stressor (e.g., more sensitive) may frequently be more critical than less sensitive pathways. However, there can be exceptions.

Time-to-effect is another intrinsic biological consideration that could be used to define critical path. For multiple paths activated at approximately the same dose, the path that could produce a relevant adverse effect earliest in the time course of exposure might be viewed as the most critical path. A simple example of time-to-effect was noted by Villeneuve et al. (2014c), who identified inhibition of glycogen synthase kinase as an molecular initiating event that could plausibly lead to impaired swim bladder inflation and subsequent reductions in young-of-year survival in fish. Although this AOP could be relevant in some exposure scenarios (e.g., exposures starting later in development), it was also acknowledged that assuming continuous exposure throughout development, impaired wnt signaling (a key event in the AOP) would be expected to cause lethal developmental abnormalities (represented as other branches in the AOP network) well before the swim bladder inflates. Thus, the swim bladder–mediated outcome was unlikely to be the critical path within that AOP network under most exposure conditions. Using a temporality layer (Knapen et al. 2018), a weighted network could be generated in which time-to-effect could be calculated for each path through the network and used to aid critical path identification (e.g., Supplemental Data, Figure S.12). Although time-to-effect information was not routinely captured in the AOP-Wiki up through 2017, updates to the guidance document (Organisation for Economic Co-operation and Development 2016a) and sections in the AOP-Wiki now prompt AOP developers to supply that information where possible. Thus, development of a temporality layer should be increasingly feasible in years to come.

It is acknowledged that some additional intrinsic and extrinsic factors may impact which path will be the critical one, or which path can shift the network toward a different critical path. For example, one might generally consider an AOP leading to complete infertility to be more critical than one that increases susceptibility to infection. However, if, under the exposure scenario being considered, an infection-driven AOP leading to acute mortality has been triggered, its intersection with the increased susceptibility pathway may make its immediate significance greater than that of long-term infertility. In other cases, intersecting AOPs in the network may enhance or strengthen the identification of a particular adverse outcome as the terminus of a critical path. For example, in the case of polypharmacological responses to synthetic glucocorticoids, immunosuppression exacerbates the effect of other AOPs, leading to decreased egg production (Margiotta-Casaluci et al. 2016; Knapen et al. 2018). Thus, information from the AOP network can actually be leveraged and utilized to support the appropriate consideration of different types of data (i.e., from endpoints or assays aligned with AOPs) based on the fact that additional critical key events may emerge in a given system of interest depending on the specific status (e.g., healthy vs diseased), environmental conditions, or polytoxicological impacts of a given stressor.

Empirical identification of critical paths. As AOP networks are developed and applied, it will become increasingly feasible to identify the most critical paths in the AOP network based on experience. As an example, we can consider the case of a chemical tested in ToxCast® that has been shown to trigger a number of different biological changes mapping to key events in an AOP network (e.g., similar to case study 3 in Knapen et al. 2018). The AOP network lays out a range of possible effects. However, after the compound is tested in vivo, it may become evident that the actual profile of observed effects followed a particular path through the network. With experience in testing more and more stressors, it may become apparent (based on the accumulated data) that among all the key events investigated, certain sequences of key events are observed more frequently than others (at least for stressors that fall within the domain tested). These essentially become empirically identified critical paths. Once these are known, other possible outcomes documented via the AOP network can be probabilistically de-emphasized, allowing one to focus on the most likely outcome(s) based on previous experience, while still recognizing other possible outcomes that could occur. The most commonly dominant paths may become those for which development of quantitative understanding and computational tools needed to infer or predict effects

along the pathway (i.e., development of quantitative AOPs; for an example, see Conolly et al. 2017) may be of higher priority. In addition, when the critical path(s) identified empirically represents a path that emerged from the network (rather than from the description of an individual AOP; Figures 2 and 5A), then identification of the critical path could plausibly lead to development of new guidelines for toxicity tests or testing strategies, in that the critical path might link together a series of endpoints and assays that had not been previously grouped as a predictive unit or motif.

Limitations of critical path identification. Although conceptually there are many benefits to the identification of critical paths, it remains to be seen whether dominant paths will really emerge through prospective AOP network-informed testing and/or retrospective AOP network-based meta-analyses. Accurate identification of critical paths will be limited by the scope of existing data. Furthermore, it is unclear whether critical paths will be conserved across species, even if the AOPs in question are relevant to their taxonomic domain, or whether certain paths will be more dominant in some species than others. As noted previously (in the *AOPs as Networks* section), the AOP knowledgebase, and thus the AOP networks, are not yet comprehensive. An undiscovered or unannotated AOP may end up being more critical within a given network than those identified through any of the methods previously described. Although there will always be a theoretical risk that the true critical path(s) may not be identified, that risk should decrease as more AOPs are described in the knowledgebase and more toxicity testing results are interpreted in an AOP network context.

Overall, the identification of critical paths within AOP networks is an emerging concept that will aid elucidation of the most fit-for-purposes assays for which there will be a good predictive value with respect to the potential adverse outcomes. By highlighting more directly the link between the measured key event and the adverse outcome, it is anticipated that this structure will ultimately allow better communication with regulators and other stakeholders and will advance the use of AOPs for risk assessment.

Interactions among AOPs

A third critical aim with regard to analysis of AOP networks is consideration of the potential interactions among AOPs. The concept of AOP interactions describes how one or more features of an AOP or its underlying biology affects another. The consequence of these interactions may be a biological outcome that is different from the one that would be observed had the interaction not occurred. As a result, the ability to understand the potential consequences of interactions among AOPs that may be activated within an organism is, arguably, one of the greatest challenges to the predictive utility of the AOP framework. It is also perhaps the most important challenge to meet given the ubiquity of exposure to multiple stressors and the fact that individual stressors may have multiple modes of action.

Interactions among AOPs can take place at many different levels of biological organization. For example, at the cellular level there can be interactions among signaling pathways, receptor cross-talk, or assembly and regulation of transcription factor complexes. Knowledge of many of these interactions has been collected and made accessible in several computational databases, such as the XTalkDB (Sam et al. 2017). Moreover, studies that describe how to leverage this knowledge to address specific biomedical challenges have been published (see Jaeger et al. 2016). Another important example of pathway interactions involves cross-talk between nuclear receptors. Numerous examples in the context of endocrine disruption have recently been reviewed by Kiyama (2016). Even when there are no directly shared key events, multiple pressures on a biological system can lead to systemic impacts such as mitochondrial energy depletion, limiting enzyme depletion, or changes in intracellular or extracellular matrix organization (Koch and Funk 2001). Likewise, as tissue or organ functions are impacted, effects in one organ can be expected to impact functions elsewhere in the body. Cross-talk between cell-signaling pathways and between various organ systems has been extensively studied over decades. Thus, it is not our intent in the present study to review all the possible biological interactions that may translate to interactions among AOPs. Instead, we aim to introduce some of the ways AOP networks can be qualitatively analyzed to gain insight into interactions and their potential consequences. A more quantitative assessment of how AOP interactions may influence the probability or severity with which downstream key events are observed is beyond the scope of the present study and is being addressed by other authors (LaLone and Hecker 2017).

Interactions and AOPs. Interactions within an AOP network can result in either an intensity or a trajectory of biological change that is different from what would be expected based on consideration of any one constituent AOP (Figure 5). Potential shifts in intensity/severity are probably the most easily identified and intuitively interpreted of the AOP interactions. Typical examples of cross-talk–induced change include additive, synergistic, and antagonistic responses (Vert and Chory 2011). These apply equally well to the potential qualitative consequences that one might expect as a result of interactions among AOPs.

From a risk assessment perspective, pathway interactions that result in additivity or synergism are of primary concern, because the interactions among multiple AOPs have the potential to amplify risk. Fortunately, convergent topologies generally associated with probable increased severity of a given adverse outcome or key event are among the easiest to qualitatively detect and interpret using AOP network analytics approaches. The ability to visually or computationally identify points at which effects of separate stressors may converge (see the *Network topology* section) to influence common downstream key events can be a basis for considering those stressors jointly in a risk assessment. This is a significant potential application of AOP networks, because currently risk assessors have relatively few tools for identifying whether potential additive or synergistic effects should be considered.

Antagonistic interactions between AOPs that would be expected to result in diminished severity are a bit more

challenging to detect based on AOP networks. This is because key events are defined as a measurable change in a biological state, expressed as an increase or decrease compared with a control/reference. Because each key event in an AOP network represents a directional change of state (see *AOPs as networks* section), counteracting effects on the same object or process will generally be defined as separate key events (separate nodes) in the network and are therefore not typically represented as convergent motifs/key events. For example, in the CYP19-AOP network (Figures 1 and 3B), estrogen receptor agonism (AOP-Wiki, Event 111; Society for the Advancement of AOPs 2017) is represented as a key event that is distinct from antagonism of the estrogen receptor (AOP-Wiki, Event 112; Society for the Advancement of AOPs 2017). This makes sense from a descriptive point of view, because both key events are functionally distinct and the downstream consequences of each are different. However, this is problematic if one wants to understand the potential antagonistic interaction between AOPs containing these key events because they have no key events in common. Instead, in this case, these parallel paths are only represented in the same AOP network diagram because the network developer identified that they represent opposite actions on the same object (the estrogen receptor) and manually added them (Figure 1). An example can be found in the T4-AOP network (Figure 2) with regard to the effects of deiodinase 3 inhibition (AOP-Wiki, Event 1153; Society for the Advancement of AOPs 2017) relative to inhibition of deiodinase 1 or 2 (AOP-Wiki, Events 1002, 1009; Society for the Advancement of AOPs 2017). In this case, the 3 AOPs appear to converge on the adverse outcome of altered amphibian metamorphosis (AOP-Wiki, Event 1101; Society for the Advancement of AOPs 2017), which, on the basis of topology alone, would suggest additivity or synergism. However, a closer look at the key events reveals that the biological effect of deiodinase 3 inhibition is the opposite of that of deiodinase 1 or deiodinase 2 inhibition, which would suggest an antagonistic interaction.

With this factor in mind, it is important that computational tools developed for AOP network extraction be able to identify key events representing opposite effects on the same object (termed same object, different action [SODA]; blue rectangles in Figures 2 and 3B). Using natural language semantics for key event titles, SODAs can be difficult to identify. However, with updates to the AOP-Wiki that now represent key events using structured 3-component ontology terms (Ives et al. 2017), it should be computationally straightforward to identify these features. The remaining questions are how to best represent them visually within the network, and how to handle them computationally when calculating network statistics or conducting quantitative analyses of AOP networks (not considered in the present study). Regardless of the final implementation, appropriate annotation of SODAs in an AOP network would have immediate use for identifying points of potential antagonistic interaction.

Perhaps the most challenging types of effects that may result from interactions among AOPs are those associated with emergent pathways (Figures 2 and 5A). Emergent pathways do not follow the trajectory of any of the individual AOPs described in the AOP-Wiki, but rather yield a new phenotypic sequence and/or outcome altogether. They may represent phenotypes or adverse outcomes that can only occur if 2 or more perturbations occur within the network. Thus, they may be difficult to define from individual stressor experiments alone. However, a significant advantage of de facto network construction using the modular AOP framework is that these more complex (and more difficult to elucidate) responses may nonetheless be coded into an AOP network as key events and key event relationships are linked to different AOPs. Emergent pathways that arise as AOP networks are constructed may provide a basis for understanding what might be otherwise unanticipated or idiosyncratic patterns of results. In this respect, one could envision scenarios in which an investigator might query the AOP knowledgebase with an observed pattern of effects (a set of key events), and then expand an AOP network around those key events to identify whether a plausible interactive mechanism, relevant to the exposure conditions of interest, exists. The emergence of novel pathways as more individual AOPs are described in the knowledgebase and linked through shared key events thus represents a unique, and potentially powerful, attribute of the AOP networks and the AOP framework as a whole.

Important considerations. As is generally the case with AOPs, the ability to understand potential interactions among AOPs and predict their consequences depends on how comprehensively the AOPs cover the relevant biology. When large data gaps exist between molecular initiating events and key events, when the key event relationships are weak, or when AOPs relevant to the modes of action of concern are sparse, it is difficult to obtain a comprehensive picture of potential interactions among AOPs or to use AOP networks to understand observed patterns of response. This should always be considered as a limiting factor when one is assessing the predictive utility of an AOP network.

On the other hand, when the pathway coverage is rich, the information can become overwhelming and may obscure rather than illuminate the critical path(s). This hearkens back to the importance of scoping and problem formulation, as well as the ability to efficiently filter the overall AOP network represented in the AOP-Wiki to find the AOPs most important to a given research or regulatory question (Knapen et al. 2018). Likewise, when AOP networks are large, topological and graph theory–based computational analyses may be critical for honing in on important network features.

Finally, consideration of interactions between AOPs that change the severity or intensity of effect on a given key event or adverse outcome leads one into the quantitative analysis of AOPs. At the level described in the present study and the companion article by Knapen et al. (2018), we focus primarily on the qualitative analysis of AOP networks. However, it is recognized that a quantitative understanding of AOPs is needed to accurately predict the probability or severity of the outcome one might expect for a given exposure scenario based on AOPs or an AOP network (Conolly et al. 2017). The development of quantitative analysis utilizing AOPs and AOP networks is

considered elsewhere (Conolly et al. 2017; LaLone and Hecker 2017).

CONCLUSIONS

The need to consider concurrent effects on multiple AOPs and their potential interactions was identified as a major theme through the SETAC Horizon Scanning exercise on "Advancing the utility of the AOP framework in research and regulation" (Lalone et al. 2017). Although this need was recognized from the inception of the AOP framework (Villeneuve et al. 2014a, 2014b), we have only recently begun to approach the critical mass of AOPs needed to start deriving AOP networks from the AOP-Wiki to analyze them and test their predictive utility. Filters and layers based on structured annotation of domains of applicability, weight of evidence, quantitative understanding of key event relationships, and other features provide the potential to derive and customize the AOP networks best suited to one's research question or problem formulation. Once derived, the AOP networks can be analyzed in a variety of ways to extract useful information. Various topological analyses can be applied to identify key features to target for assay or model development. Critical paths based on risk assessment goals, biological attributes, or empirical testing can be defined. Together, these approaches, along with tools still under development, can help us to identify and understand the complex interactions that may occur when multiple AOPs are activated in different contexts. Broadly speaking, all these approaches are in their infancy with regard to understanding and illustrating their practical utility and limitations. The concepts described serve as a starting point to aid the ongoing development of the AOP-Wiki and associated software applications like AOPXplorer (Burgoon 2017). In addition, the ideas we present can inform the design of applications case studies that will put these concepts to the test. Along with other articles associated with the April 2017 SETAC Pellston Workshop on Advancing the Adverse Outcome Pathway Framework (LaLone and Hecker 2017), the present study serves the ongoing development of the AOP framework as a critical concept to support 21st century approaches to toxicological research and regulation.

Supplemental Data—The Supplemental Data are available on the Wiley Online Library at DOI: 10.1002/etc.4124.

Acknowledgment—The authors thank the Society of Environmental Toxicology and Chemistry (SETAC) Pellston Workshop co-chairs, C. LaLone and M. Hecker, for their coordination, organization, and guidance of the workshop. We acknowledge the other workshop participants (see LaLone and Hecker 2017) for their stimulating discussions and feedback and the respondents to the Horizon Scanning efforts for the charge questions and themes that informed our discussion. We gratefully acknowledge the SETAC North America staff, in particular G. Schiefer, N. Mayo, and T. Schlekat, who provided support to the workshop co-chairs, steering committee, and workshop participants before, during, and after the Pellston Workshop. Funding for the workshop was provided by SETAC, the US Environmental Protection Agency, the American Cleaning Institute, the European Chemical Industry Council Long-Range Research Initiative, Chevron Environmental Management, the European Center for Ecotoxicology and Toxicology of Chemicals, the European Commission Joint Research Centre, the European Crop Protection Association, ExxonMobil, Humane Society International, The Humane Society of the United States, the Human Toxicology Project Consortium, Syngenta, and Unilever. In addition, we thank the groups from academia, industry, and government who supported the participants' travel. Finally, we thank S. Edwards for review and input on the draft manuscript.

Disclaimer—The contents of the present study represent the personal opinions of the authors and neither constitute, nor necessarily reflect, the policies or viewpoints of their employers or institutes.

Data Availability—All data reported in the present study are available via the Supplemental Data files or from the Society for the Advancement of Adverse Outcome Pathways (2017).

REFERENCES

Alon U. 2007. Network motifs: Theory and experimental approaches. *Nat Rev Genet* 8:450–461.

Becker RA, Ankley GT, Edwards SW, Kennedy SW, Linkov I, Meek B, Sachana M, Segner H, Van Der Burg B, Villeneuve DL, Watanabe H, Barton-Maclaren TS. 2015. Increasing scientific confidence in adverse outcome pathways: Application of tailored Bradford-Hill considerations for evaluating weight of evidence. *Regul Toxicol Pharmacol* 72:514–537.

Burgoon L. 2017. Cytoscape App Store, AOPXplorer. [cited 2017 August 7]. Available from: http://apps.cytoscape.org/search?q.aopxplorer

Caldarelli G, Catanzaro M. 2012. *Networks: A Very Short Introduction*. Oxford University, Oxford, UK.

Conolly RB, Ankley GT, Cheng W, Mayo ML, Miller DH, Perkins EJ, Villeneuve DL, Watanabe KH. 2017. Quantitative adverse outcome pathways and their application to predictive toxicology. *Environ Sci Technol* 51:4661–4672.

Cytoscape. 2017. NetworkAnalyzer Online Help. [cited 2017 July 28]. Available from: http://med.bioinf.mpi-inf.mpg.de/netanalyzer/help/2.7/index.html

Friedlander T, Mayo AE, Tlusty T, Alon U. 2015. Evolution of bow-tie architectures in biology. *PLoS Comput Biol* 11:e1004055.

Huber W, Carey VJ, Long L, Falcon S, Gentleman R. 2007. Graphs in molecular biology. *BMC Bioinformatics* 8(Suppl 6):S8.

Ives C, Campia I, Wang R-L, Wittwehr C, Edwards S. 2017. Creating a structured AOP knowledgebase via ontology-based annotations. *Appl In Vitro Toxicol* 3:298–311.

Jaeger S, Igea A, Arroyo R, Alcade V, Canovas B, Orozco M, Nebreda AR, Aloy P. 2016. Quantification of pathway cross-talk reveals novel synergistic drug combinations for breast cancer. *Cancer Res* 77:1–11.

Kitsak M, Havlin S, Paul G, Riccaboni M, Pammolli F, Stanley HE. 2007. Betweenness centrality of fractal and nonfractal scale-free model networks and tests on real networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 75:056115.

Kiyama R. 2016. Endocrine disruptor actions through receptor crosstalk. *Environ Biotechnol* 12:1–16.

Knapen D, Angrish MM, Fortin MC, Katsiadaki I, Leonard M, Margiotta-Casaluci L, Munn S, O'Brien JM, Pollesch NL, Smith LC, Zhang X, Villeneuve DL. 2018. Adverse outcome pathway networks. I: Development and applications. *Environ Toxicol Chem* 37:1723–1733, (*this issue*).

Koch T, Funk RH. 2001. Cellular dysfunction in the pathogenesis of organ failure. New insights from molecular and cell biology. *Anaesthesist* 50:742–749 (in German).

LaLone CA, Hecker M, co-chairs. 2017. Society of Environmental Toxicology and Chemistry (SETAC) Pellston Workshop: Advancing the Adverse Outcome Pathway Concept: An International Horizon Scanning Approach, April 2-6, 2017, Cornwall, ON, Canada. Society for the Advancement of AOPs, Durham, North Carolina, USA. [cited 2017 August 8]. Available from: http://www.saaop.org/workshops/pellston2017.html

LaLone CA, Ankley GT, Belanger SE, Embry MR, Hodges G, Knapen D, Munn S, Perkins EJ, Rudd MA, Villeneuve DL, Whelan M, Willett C, Zhang X, Hecker M. 2017. Advancing the adverse outcome pathway framework—An international Horizon Scanning approach. *Environ Toxicol Chem* 36:1411–1421.

Lewis TG. 2009. *Network Science: Theory and Applications*. John Wiley & Sons, Hoboken, NJ, USA.

Margiotta-Casaluci L, Owen SF, Huerta B, Rodriguez-Mozaz S, Kugathas S, Barceló D, Rand-Weaver M, Sumpter JP. 2016. Internal exposure dynamics drive the adverse outcome pathways of synthetic glucocorticoids in fish. *Sci Rep* 6:srep21978.

Milo R, Shen-Orr S, Itzkovitz S, Kashtan N, Chklovskii D, Alon U. 2002. Network motifs: Simple building blocks of complex networks. *Science* 298:824–827.

Netzwerkerin. Eccentricity and closeness. 2017. [cited 2017 July 28]. Available from: https://sites.google.com/site/netzwerkerin/home/closeness

Newman MEJ. 2003. The structure and function of complex networks. *SIAM Rev* 45:167–256.

Organisation for Economic Co-operation and Development. 2016a. User's handbook supplement to the guidance document for developing and assessing adverse outcome pathways. Paris, France. [cited 2017 August]. Available from: https://one.oecd.org/document/ENV/JM/MONO(2016)12/en/pdf

Organisation for Economic Co-operation and Development. 2016b. Guidance document for the use of adverse outcome pathways in developing integrated approaches to testing and assessment (IATA). Series on Testing and Assessment, No. 260. ENV/JM/MONO(2016)67. Paris, France.

Pavlopoulos GA, Secrier M, Moschopoulos CN, Soldatos TG, Kossida S, Aerts J, Schneider R, Bagos PG. 2011. Using graph theory to analyze biological networks. *BioData Mining* 4:10.

Sam SA, Teel J, Tegge AN, Bharadwaj A, Murali TM. 2017. XTalkDB: A database of signaling pathway crosstalk. *Nucleic Acids Res* 45: D432–D439.

Skiena S. 1990. Topological sorting, §5.4.3. In *Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*, 1st ed. Basic Books, New York, NY, USA, pp 208–209.

Society for the Advancement of Adverse Outcome Pathways. 2017. AOP-Wiki. [cited 2017 August 8]. Available from: http://aopwiki.org

Trudeau RJ. 2013. *Introduction to Graph Theory*. Dover Publications, New York, NY, USA.

US Environmental Protection Agency, Risk Assessment Forum. 1992. Framework for ecological risk assessment. EPA/630/R-92/001.Washington, DC.

US Environmental Protection Agency, Risk Assessment Forum. 1998. Guidelines for ecological risk assessment. EPA/630/R-95/002F. Washington, DC.

US Environmental Protection Agency, Risk Assessment Forum. 2014. Framework for human health risk assessment to inform decision making. EPA/100/R-14/001.Washington DC.

Vazquez A, Dobrin R, Sergi D, Eckmann JP, Oltvai ZN, Barabasi AL. 2004. The topological relationship between the large-scale attributes and local interaction patterns of complex networks. *Proc Natl Acad Sci U S A* 101:17940–17945.

Vert G, Chory J. 2011. Crosstalk in cellular signaling: Background noise or the real thing? *Dev Cell* 21:985–991.

Villeneuve DL. 2017. AOP-Wiki. Society for the Advancement of Adverse Outcome Pathways. [cited 2017 July 18]. Available from: https://aopwiki.org/aops/25

Villeneuve DL, Crump D, Garcia-Reyero N, Hecker M, Hutchinson TH, LaLone CA, Landesmann B, Lettieri T, Munn S, Nepelska M, Ottinger MA, Vergauwen L, Whelan M. 2014a. Adverse outcome pathway (AOP) development. I: Strategies and principles. *Toxicol Sci* 142:312–320.

Villeneuve DL, Crump D, Garcia-Reyero N, Hecker M, Hutchinson TH, LaLone CA, Landesmann B, Lettieri T, Munn S, Nepelska M, Ottinger MA, Vergauwen L, Whelan M. 2014b. Adverse outcome pathway development. II: Best practices. *Toxicol Sci* 142:321–330.

Villeneuve D, Volz DC, Embry MR, Ankley GT, Belanger SE, Leonard M, Schirmer K, Tanguay R, Truong L, Wehmas L. 2014c. Investigating alternatives to the fish early-life stage test: A strategy for discovering and annotating adverse outcome pathways for early fish development. *Environ Toxicol Chem* 33:158–169.

Weisstein EW. 2017. Topological sort. In MathWorld—A Wolfram Web Resource. [cited 2017 August 8]. Available from: http://mathworld.wolfram.com/TopologicalSort.html

Worth AP, Patlewicz G. 2016. Integrated approaches to testing and assessment. *Adv Exp Med Biol* 856:317–342.

